Oracle® Fusion Middleware

Exalogic Enterprise Deployment Guide

Release EL X2-2 and EL X3-2

E18479-10

October 2012

ORACLE®

Oracle Fusion Middleware Exalogic Enterprise Deployment Guide, Release EL X2-2 and EL X3-2

E18479-10

Primary Author: Gururaj

Contributing Author: Showvik Roychowdhuri

Contributors: Ballav Bihani, James Bayer, Kaizhe Huang, Michael Lehmann, Naresh Revanuru, Ola Tørudbakken, Peter Bower, Rich Mousseau, Sridhar Ranganathan, and Will Lyons.

Contents

3 Network, Storage, and Database Preconfiguration

4 Installing Oracle Software

5 Configuring Oracle Fusion Middleware

6 Configuring Oracle HTTP Server

7 Enabling Exalogic-Specific Enhancements in Oracle WebLogic Server 11*g* Release 1 (10.3.4)

8 Deploying a Sample Web Application to an Oracle WebLogic Cluster

9 Managing the Topology

Preface

This preface describes the audience, contents and conventions used in the Oracle Fusion Middleware Enterprise Deployment Guide for Oracle Exalogic.

Intended Audience

This guide is intended for administrators who are responsible for installing and configuring Oracle Exalogic enterprise deployments.

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs.

Access to Oracle Support

Oracle customers have access to electronic support through My Oracle Support. For information, visit http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info or visit http://www.oracle.com/pls/topic/lookup?ctx=accid=trs if you are hearing impaired.

Related Documents

For more information, see the following documents:

- *Oracle Fusion Middleware Administrator's Guide*
- *Oracle Exalogic Machine Owner's Guide*
- *Oracle Exalogic Release Notes*
- *Oracle Fusion Middleware 11 g Release 1 (11.1.1) Documentation*

Conventions

The following text conventions are used in this document:

Convention	Meaning
boldface	Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.

Convention	Meaning
italic	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.
`monospace`	Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter.

Before You Begin

This document provides the basic procedures for installing and configuring Oracle WebLogic Server 10.3.4 and Oracle Coherence 3.6.1 on Oracle Exalogic. It applies only to the Exalogic X2-2 hardware platform, and it does not apply to the new Exalogic X3-2.

For information about deploying other versions of Oracle WebLogic Server and Oracle Coherence on Exalogic, or for information about deploying these products on newer releases of the Exalogic hardware, you can use the information in this document as a general framework; the specifics may vary.

This document may contain some information about Oracle Exalogic-specific features and optimizations in Oracle Fusion Middleware products. For up-to-date and detailed information about these optimizations, see the Oracle Fusion Middleware documentation.

1 Enterprise Deployment Overview

This chapter introduces enterprise deployment reference topologies and configuration scenario for Oracle Exalogic. It contains the following sections:

1.1 What is Enterprise Deployment?

Enterprise deployment is an Oracle best practices blueprint based on proven Oracle high-availability and security technologies and recommendations for Oracle Exalogic. The best practices described in these blueprints span all Oracle products across the entire technology stack: Oracle Database, Oracle Fusion Middleware, Oracle Exalogic machine, and Oracle Enterprise Manager Grid Control.

An Oracle Fusion Middleware enterprise deployment:

- considers various business service level agreements (SLA) to make high-availability best practices as widely applicable as possible
- leverages database grid servers and storage grid with low-cost storage to provide highly resilient, lower cost infrastructure
- uses results from extensive performance impact studies for different configurations to ensure that the high-availability architecture is optimally configured to perform and scale to business needs
- enables control over the length of time to recover from an outage and the amount of acceptable data loss from a natural disaster
- uses Oracle best practices and recommended architecture.

For more information about high availability practices, go to `http://www.oracle.com/technology/deploy/availability/htdocs/maa.htm`.

1.2 Prerequisities

Setup and commissioning of Oracle Exalogic machine, including initial storage and networking configuration, as described in *Oracle Fusion Middleware Exalogic Machine Owner's Guide*.

1.3 Terminology

This section provides information about Oracle Fusion Middleware concepts and terminologies that are related to administering Oracle Fusion Middleware.

- **Middleware home:** A Middleware home consists of the Oracle WebLogic Server home, and, optionally, one or more Oracle homes.
- **Oracle home**: An Oracle home contains installed files necessary to host a specific product. An Oracle home resides within the directory structure of the Middleware home. Each Oracle home can be associated with multiple Oracle instances or Oracle WebLogic Server domains.
- **WebLogic Server home**: A WebLogic Server home contains installed files necessary to host a WebLogic Server. The WebLogic Server home directory is a peer of Oracle home directories and resides within the directory structure of the Middleware home.
- **Oracle WebLogic Server Domain**: An Oracle WebLogic Server administration domain is a logically related group of Java components. A domain includes a special WebLogic Server instance called the Administration Server, which is the central point from which you configure and manage all resources in the domain. Usually, you configure a domain to include additional WebLogic Server instances called Managed Servers. You deploy Java components, such as Web applications, EJBs, and Web services, and other resources, to the Managed Servers and use the Administration Server for configuration and management purposes only.

  The directory structure of a domain is separate from the directory structure of the WebLogic Server home.
- **failover**: When a member of a high availability system fails unexpectedly (unplanned downtime), in order to continue offering services to its consumers, the system undergoes a failover operation. If the system is an active-passive system, the passive member is activated during the failover operation and consumers are directed to it instead of the failed member. The failover process can be performed manually, or it can be automated. If the system is an active-active system, the failover is performed by the load balancer entity serving requests to the active members. If an active member fails, the load balancer detects the failure and automatically redirects requests for the failed member to the surviving active members. See *Oracle Fusion Middleware High Availability Guide* for information on active-active and active-passive systems.
- **failback**: After a system undergoes a successful failover operation, the original failed member can be repaired over time and be re-introduced into the system. If desired, a failback process can be initiated to activate this member and deactivate the other. This process reverts the system to its pre-failure configuration.
- **server migration**: A feature of WebLogic Server required for applications that have critical data, such as persistent JMS or transaction logs. If a Managed Server hosting the application fails, the server migration feature of WebLogic Server ensures that the application availability is not affected. This task is not required for applications that do not have persistent JMS or transaction logs.

- **shared storage**: Shared storage refers to the Sun ZFS Storage 7320 appliance that is accessible by all compute nodes in the Oracle Exalogic Machine. All compute nodes in the Exalogic machine can access this storage appliance simultaneously for both read and write operations.

  Among other things, the following artifacts are located on the Sun ZFS Storage 7320 appliance:

  – Middleware Home software

  – Oracle WebLogic Server domains

  – Oracle WebLogic Server log files

  – JMS persistence logs

  – JTA logs (where applicable)

  – Application-specific artifacts, such as data, images, and so on

  – Oracle Linux (OL) or Oracle Solaris operating system crash dumps, patches, and syslogs

  Note: The factory setting for Oracle Exalogic machine is to store syslogs on the local storage of compute nodes. However, you can configure log rotation to store syslogs on the Sun ZFS Storage 7320 appliance, as necessary.

- **compute node**: A physical machine in Exalogic rack that is meant for running an Oracle Fusion Middleware instance at any given time and has been configured to have a backup compute node.
- **network host name**: Network host name is a name assigned to an IP address either through the `/etc/hosts` file or through DNS resolution. This name is visible in the network that the machine to which it refers to is connected. Often, the network host name and physical host name are identical. However, each machine has only one physical host name but may have multiple network host names. Thus, a machine's network host name may not always be its physical host name.
- **physical host name**: This guide differentiates between the terms physical host name and network host name. This guide uses physical host name to refer to the "internal name" of the current machine. On UNIX, this is the name returned by the `hostname` command.

  Physical host name is used by Oracle Fusion Middleware to reference the local host. During installation, the installer automatically retrieves the physical host name from the current machine and stores it in the Oracle Fusion Middleware configuration metadata on disk.
- **physical IP**: Physical IP refers to the IP of a machine on the network. In almost all cases, it is normally associated with the physical host name of the machine (see the definition of the physical host name). In contrast to a virtual IP or floating IP, it is always associated with the same machine when on a network.
- **switchover**: During normal operation, active members of a system may require maintenance or upgrading. A switchover process can be initiated to allow a substitute member to take over the workload performed by the member that requires maintenance or upgrading, which undergoes planned downtime. The switchover operation ensures continued service to consumers of the system.

- **switchback**: When a switchover operation is performed, a member of the system is deactivated for maintenance or upgrade. When the maintenance or upgrade is completed, the system can undergo a switchback operation to activate the upgraded member and bring the system back to the pre-switchover configuration.
- **virtual host name**: Virtual host name is a network addressable host name that maps to one or more physical machines via a load balancer. For load balancers, the name "virtual server name" is used interchangeably with virtual host name in this book. A load balancer can hold a virtual host name on behalf of a set of servers, and clients communicate indirectly with the machines using the virtual host name.

  Note: Whenever the term "virtual host name" is used in this document, it is assumed to be associated with a virtual IP address. In cases where just the IP address is needed or used, it will be explicitly stated.

- **virtual IP**: Generally, a virtual IP can be assigned to a load balancer. To present a single system view of a cluster to network clients, a virtual IP serves as an entry point IP address to the group of servers which are members of the cluster. A virtual IP can be assigned to a server load balancer.

  A load balancer uses a virtual IP as the entry point to a set of servers. These servers tend to be active at the same time. This virtual IP address is not assigned to any individual server but to the load balancer which acts as a proxy between servers and their clients.
- **floating IP**: Floating IP is the IP assigned to one of the WebLogic Managed Servers in a Weblogic cluster to allow for server migration.

1.4 Benefits of Oracle Recommendations

The Oracle Fusion Middleware configurations discussed in this guide are designed to ensure security of all invocations, maximize hardware resources, and provide a reliable, standards-compliant system for enterprise computing with a variety of applications.

- Section 1.4.1, "Built-in Security"
- Section 1.4.2, "High Availability"
- Section 1.4.3, "Performance"
- Section 1.4.4, "Application Isolation"

The security and high availability benefits of the Oracle Fusion Middleware configurations are realized through isolation in firewall zones and replication of software components.

1.4.1 Built-in Security

The Enterprise Deployment architectures are secure because every functional group of software components is isolated in its own Demilitarized Zone (DMZ), and all traffic is restricted by protocol and port. The following characteristics ensure security at all needed levels, as well as a high level of standards compliance:

- Configure external load balancers to redirect all external communication received on port 80 to port 443.

Note: The Oracle Technology Network (http://www.oracle.com/technology/index.html) provides a list of validated load balancers and their configuration at http://www.oracle.com/technology/products/ias/hi_av/Tested_LBR_FW_SSLAccel.html.

- Communication from external clients does not go beyond the Load Balancing Router level.
- No direct communication from the Load Balancing Router to the data tier is allowed.
- Components are separated in different protection zones: the web tier, application tier, and the data tier.
- Direct communication between two firewalls at any one time is prohibited.
- If communication begins in one firewall zone, it must end in the next firewall zone.
- All communication between components across protection zones is restricted by port and protocol, according to firewall rules.

1.4.2 High Availability

The enterprise deployment architectures are highly available, because each component or functional group of software components is replicated on a different computer, and configured for component-level high availability.

1.4.3 Performance

Oracle Exalogic uses InfiniBand as the I/O fabric technology. InifiniBand provides a high throughput, low latency, and scalable fabric that is suitable for fabric consolidation of inter-processor communication, network and Storage. It is optimized for cluster and storage traffic.

Regardless of the design of the application, Oracle Exalogic offers a multitude of capabilities that dramatically improve the overall performance and reliability of the application. To benefit from the features and capabilities of Oracle Exalogic, Oracle WebLogic Suite 11*g* users only need to deploy their applications to the Exalogic machine; no code changes or rearchitecture of applications is necessary.

1.4.4 Application Isolation

Oracle Exalogic provides a high degree of isolation among concurrently deployed applications that have diverse security, reliability, and performance requirements. It creates a default IP over InfiniBand (IPoIB) link and an Ethernet over InfiniBand (EoIB) interface during initial configuration. All compute nodes in the Exalogic Machine are members of the default InfiniBand partition.

The most common model for application isolation involves multiple IP subnetting, in which the most mission-critical applications are assigned their own IP subnets layered above the default IPoIB link. In this model, some subnets may also contain applications that have less stringent or otherwise different resource requirements. Other subnets may host WebLogic domains, which contain multiple applications, such as those dedicated to a given department or line of business, or even used for application testing and development.

1.5 Overview of Oracle Exalogic Configured Environment

Before you start implementing the Oracle Exalogic enterprise deployment topology, you should understand the current state of the Exalogic environment.

It is assumed that you have completed all tasks described in the *Oracle Fusion Middleware Exalogic Machine Owner's Guide*, which discusses your data center site preparation, Oracle Exalogic machine commissioning, initial networking configuration including IP address assignments, and initial setup of the Sun ZFS Storage 7320 appliance.

This section describes the state of the Exalogic configured environment before enterprise deployment.

It discusses the following topics:

- Section 1.5.1, "Network"
- Section 1.5.2, "Sun ZFS Storage 7320 appliance"
- Section 1.5.3, "Oracle Software"

1.5.1 Network

Before you start configuring the enterprise deployment topology, you must run the `Exalogic Configuration Utility` to complete the following tasks, as described in the chapter "Initial Configuration of Exalogic Machine Using Oracle Exalogic Configuration Utility" in the *Oracle Fusion Middleware Exalogic Machine Owner's Guide*:

- Configuration of IP addresses for all Exalogic compute nodes and the Sun ZFS Storage 7320 appliance.
- Configuration of InfiniBand gateway switches.
- Configuration of the Cisco Ethernet management switch.
- Setup and verification of the default IP over InfiniBand (IPoIB) link spanning all compute nodes.
- Setup and verification of the default Ethernet over InfiniBand (EoIB) link for connectivity with components of the topology running on Ethernet.
- Configuration of the default InfiniBand partition that covers all of the compute nodes in Exalogic Machine.

1.5.2 Sun ZFS Storage 7320 appliance

The initial configuration of the Sun ZFS Storage 7320 appliance in your Oracle Exalogic machine is completed at the time of manufacturing. For more information about default shares (Exported File Systems), see the "Default Storage Configuration" section in the *Oracle Fusion Middleware Exalogic Machine Owner's Guide*.

After completing this initial configuration, you must proceed to create custom shares, as described in Section 3.4.2, "Setting Up Enterprise Deployment Storage Configuration".

1.5.3 Oracle Software

Oracle Linux 5.5 or Oracle Solaris 11 is pre-installed on each of the compute nodes in your Oracle Exalogic machine.

You must download the Oracle WebLogic 10.3.4 software installer and run the installation program on one of the compute nodes. You must save the installation

binaries, including Oracle Middleware Home, on a shared file system on the Sun ZFS Storage 7320 appliance. Before you can do so, you must configure shared storage by creating a Project and defining shares and LUNs to set up the directory structure, as necessary. Note down the mount point for such shares, so you can mount the required locations or directories from Exalogic compute nodes.

For more information, see Chapter 3, "Network, Storage, and Database Preconfiguration" and Chapter 4, "Installing Oracle Software".

Note: You can download the Oracle WebLogic 10.3.4 software from http://edelivery.oracle.com. Select **Oracle Fusion Middleware** as the Product Pack, **Linux x86-64** or **Oracle Solaris on x86-64 (64-bit)** as the Platform, and **Oracle Fusion Middleware 11g Media Pack for Exalogic** as the Media Pack.

1.6 Administrator Roles and Permissions

Administration and management of Oracle Exalogic may span multiple specialized roles and separate departments in organizations. This is due to the integrated nature of Exalogic that combines multiple compute servers, shared storage and shared networking infrastructure. Oracle recommends that you align the planned use of Exalogic to the appropriate roles in your organization. For example, during initial deployment and day-to-day operations, you may consider roles, such as the following:

- **Machine Administrator** - Administers all resources internal to Exalogic. This is the only role with `root` credentials on compute nodes.
- **Storage Administrator** - Administers the Sun ZFS Storage 7320 appliance.
- **Network Administrator** - Administers the InfiniBand gateway switches and management switches in the Oracle Exalogic machine. This administrator may also have permissions to configure resources external to Exalogic, such as hardware load balancers, firewalls, and web servers in the Web Tier.
- **Database Administrator** - Administers database connectivity from software running in Oracle Exalogic.
- **Department Administrator** - Administers X4170 M2 compute nodes in the Oracle Exalogic machine as non-root account, such as user `weblogic` in the operating system group `oracle` that has permissions to install, deploy, configure, and manage department processes and resources.
- **WebLogic Domain Administrator** - Administers a department's middleware, such as WebLogic Server domains and Node Manager. This user likely has more restricted operating system privileges than the Department Administrator.
- **Operations and Management Administrator** - A user that does not have rights to deploy or manipulate running applications, but is able to access management tools, such as Enterprise Manager Grid Control for monitoring purposes.

1.7 Task Roadmap

Table 1–1 lists high-level enterprise deployment tasks for Oracle Linux or Solaris physical environments.

Table 1–1 Enterprise Deployment Tasks for Oracle Linux or Solaris Physical Environments

Step	Description	For More Information
1	Familiarize yourself with Exalogic reference topologies.	See Enterprise Deployment Overview and Reference Topology and Slicing Diagram
2	Examine your Exalogic machine rack and compute nodes. Ensure that the base operating system (Oracle Linux or Solaris 11 Express) is installed on the compute nodes, and the primary IPoIB and EoIB interfaces are set up.	See Overview of Oracle Exalogic Configured Environment
3	Review and understand the horizontal slicing of an Exalogic machine and the example configuration scenario. Slicing of an Exalogic machine quarter rack is included as an example.	See Example: Horizontal Slicing Within Exalogic Machine Quarter Rack and Example Configuration Scenario for Exalogic x86 Physical Machines
4	Prepare your network, database, and storage for enterprise deployment.	See Network, Storage, and Database Preconfiguration
5	Install Oracle software.	See Installing Oracle Software
6	Configure Oracle Fusion Middleware software.	See Configuring Oracle Fusion Middleware
7	Configure Oracle HTTP Server.	See Mandatory: Configuring Oracle HTTP Server for Administration Server and Managed Servers
8	Enable Exalogic-specific optimizations in WebLogic Server.	See Enabling Exalogic-Specific Enhancements in Oracle WebLogic Server 11g Release 1 (10.3.4)
9	Deploy a sample application to the WebLogic cluster that is configured to run on Exalogic compute nodes.	See Deploying a Sample Web Application to an Oracle WebLogic Cluster
10	Manage the enterprise deployment topology.	See Managing the Topology
11	Monitor the software in the enterprise deployment topology by using Oracle Enterprise Manager Grid Control.	See Monitoring the Topology Using Oracle Enterprise Manager Grid Control

2

Reference Topology and Slicing Diagram

This chapter describes Exalogic enterprise deployment reference topologies. The instructions and diagrams in this guide describe two scenarios, to which variations may be applied. The chapter also includes the horizontal slicing diagram for an example Exalogic machine configuration.

It contains the following sections:

2.1 Scenario1: Exalogic Machine Connected to Oracle Database or RAC over 10 Gb Ethernet

Figure 2–1 illustrates how an Exalogic machine is connected to an Oracle database or RAC over 10 Gb Ethernet in the enterprise deployment reference topology.

Figure 2–1 Exalogic Enterprise Deployment Reference Topology with Oracle Database over Ethernet in the Data Tier

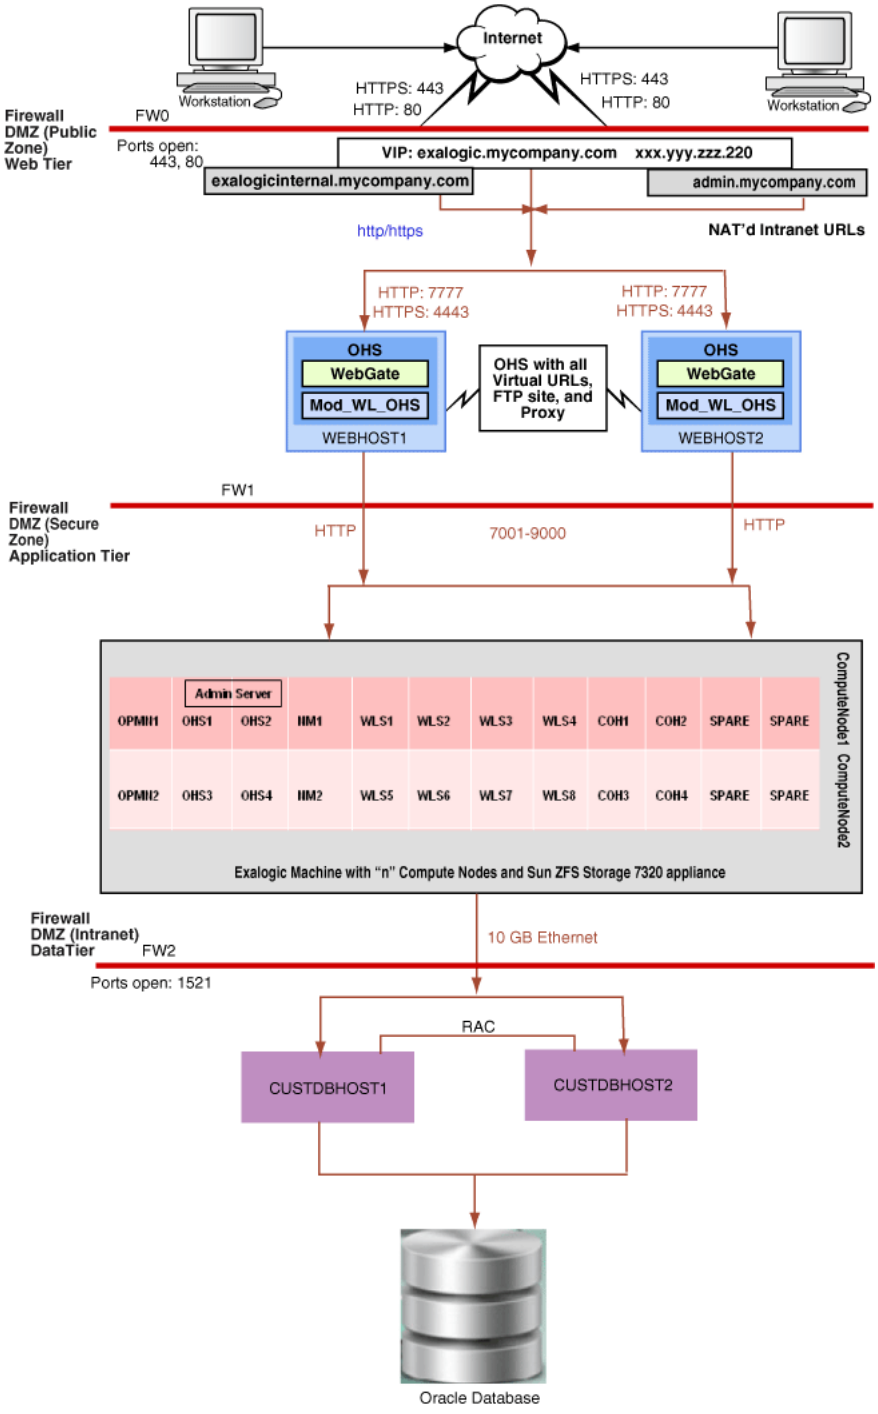

You can use the procedures described in this guide to configure this reference topology. In addition, you must create two vNICs (`vnic0/vnic1`) for connecting an Oracle Exalogic machine to Oracle database or RAC over a 10 GB Ethernet link. For information about creating these vNICs, see Section 3.5.2, "Connecting to Oracle Database Over Ethernet".

Note: Although the figure shows 7001-9000 ports, ensure that only ports that are used will be opened.

2.2 Scenario2: Exalogic Machine Connected to Oracle Exadata Database Machine via InfiniBand

Figure 2–2 illustrates how an Exalogic machine is connected to an Oracle Exadata Database Machine together on the same InfiniBand fabric.

Figure 2–2 Exalogic Enterprise Deployment Reference Topology with Oracle Exadata Database Machine

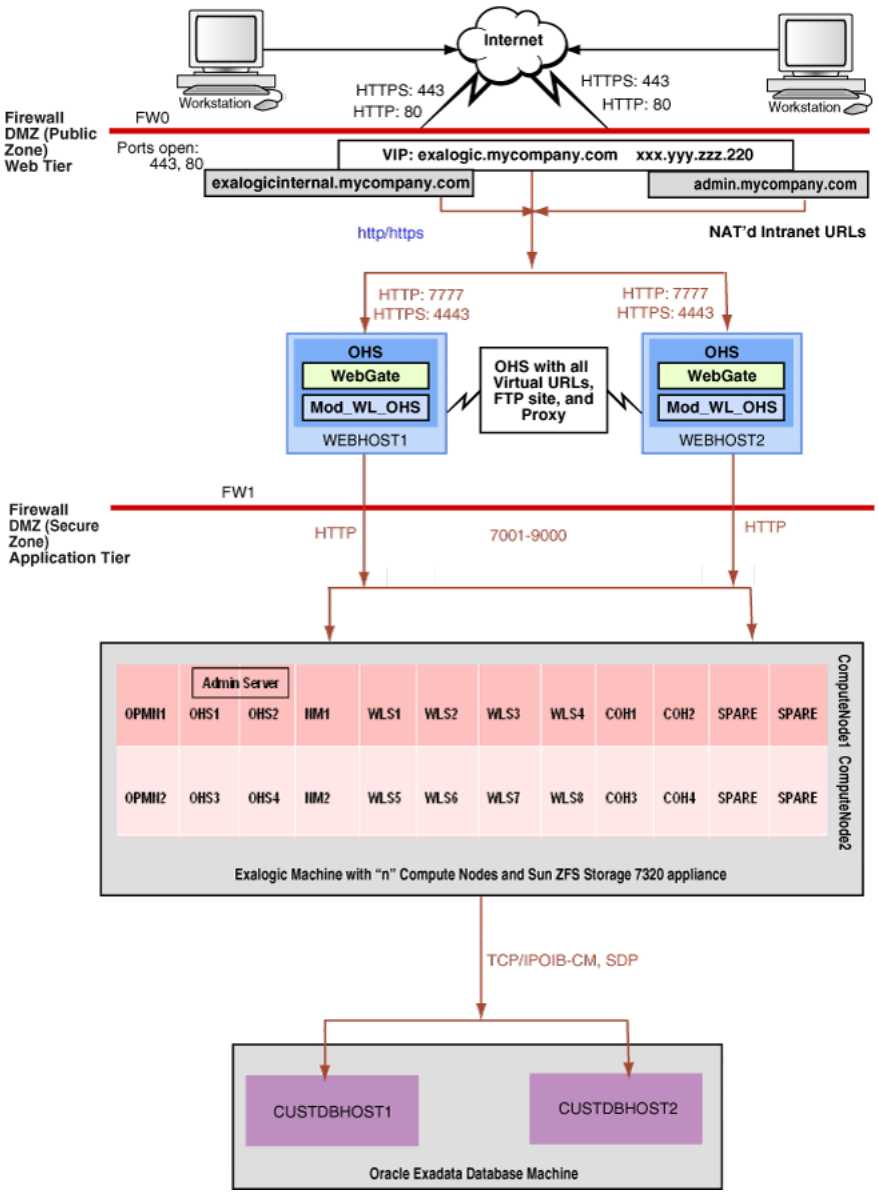

You can use the procedures described in this guide to configure this reference topology. For information about connecting an Oracle Exalogic machine to Oracle

Exadata Database Machine, see the *Oracle Fusion Middleware Exalogic Machine Multirack Cabling Guide*.

Note: Although the figure shows 7001-9000 ports, ensure that only ports that are used will be opened.

2.3 Processor Cores for Exalogic X86 Machines

The number of processor cores on an Exalogic compute node depends on the machine version. For example, a compute node on a standard Exalogic X2-2 machine has two 6-core processors (12 cores in total), and a compute node on a standard X3-2 machine has two 8-core processors (16 cores in total).

Although the reference topologies show two compute nodes, the number of compute nodes available for your application deployment and configuration depends on your Exalogic machine configuration. However, you can extrapolate and use the example scenario described in this guide to set up an enterprise reference topology based on your specific requirements.

2.4 Introduction to Tiers

This section introduces the following tiers:

- Web Tier
- Application Tier
- Data Tier

2.4.1 Web Tier

Nodes in the web tier are located in the DMZ public zone.

The configuration example described in this guide uses Oracle HTTP Server as the web server.

In this tier, two nodes `WEBHOST1` and `WEBHOST2` run Oracle HTTP Server configured with `mod_wl_ohs`.

Through `mod_wl_ohs`, which allows requests to be proxied from Oracle HTTP Server to WebLogic Server, Oracle HTTP Server forwards the requests to WebLogic Server running in the application tier.

The web tier also includes an external load balancer router to handle external requests. External requests are sent to the virtual host names configured on the load balancer. The load balancer then forwards the requests to Oracle HTTP Server.

On the firewall protecting the web tier, only the HTTP ports are open: 443 for HTTPS and 80 for HTTP.

Note: A separate installation of Oracle HTTP Server inside the Exalogic machine in the application tier provides internal load balancing and routes requests via IP over InfiniBand (IPoIB) and Ethernet over InfiniBand (EoIB). This installation is optional.

2.4.1.1 Parameters for Web Server Plug-Ins

You enter the parameters for each Web server plug-in special configuration files. Each Web server has a different name for this configuration file and different rules for formatting the file. In this guide, you must set the Oracle WebLogic cluster parameter to specify the floating IP addresses of Oracle WebLogic Managed Server and their ports in the configuration file used by the plug-in for Oracle HTTP Server.

For more information, see the *Oracle Fusion Middleware Using Web Server 1.1 Plug-Ins with Oracle WebLogic Server*. In addition, see Chapter 6, "Configuring Oracle HTTP Server".

2.4.2 Application Tier

Nodes in the application tier are located in the DMZ secure zone.

In this tier, Oracle WebLogic Server is configured with Managed Servers for running Oracle Fusion Middleware components on compute nodes in Oracle Exalogic machine. The number of compute nodes in Oracle Exalogic machine depends on your purchased hardware configuration, such as Oracle Exalogic machine full rack (30 compute nodes), Oracle Exalogic machine half rack (16 compute nodes), and Oracle Exalogic Machine quarter rack (8 compute nodes).

Optionally, you can configure and run Oracle HTTP Server instances on your Exalogic compute nodes if you wish to load balance traffic between WebLogic server instances running on the compute nodes. Additionally, Oracle HTTP Server instances running on Exalogic compute nodes provide routing of requests from Oracle HTTP Server to WebLogic server instances via IPoIB and EoIB.

Note: The topology diagrams show only two Oracle Exalogic compute nodes (`ComputeNode1` and `ComputeNode2`) in the application tier. The configuration procedures described in this guide are based on horizontal slicing (aaaaaaa) within a single Oracle Exalogic machine. In addition, it considers a simple configuration scenario, including two clusters of 8 Managed Servers each on `ComputeNode1` and `ComputeNode2`, Node Manager running on both of the compute nodes, and the Administration Server running on `ComputeNode1`. For the Administration Server running on `ComputeNode1`, the second compute node `ComputeNode2` is used as the back-up machine.

You can use this example configuration to identify and implement your specific configuration requirements.

2.4.3 Data Tier

Nodes in the data tier are located in the most secured network zone (the intranet).

In the data tier, you should consider the following configuration scenarios:

- Oracle Exalogic machine connected to Oracle database or RAC over 10 Gb Ethernet
- Oracle Exalogic machine connected to Oracle Exadata Database Machine over InfiniBand

If you are connecting your Oracle Exalogic machine to an Oracle database or RAC over Ethernet, you must create vNICS and VLANs, as necessary. For more information, see Section 3.5.2, "Connecting to Oracle Database Over Ethernet".

If you are connecting your Oracle Exalogic machine to Oracle Exadata Database Machine to communicate via InfiniBand, you must connect the Oracle Exalogic machine to Oracle Exadata Database Machine, as described in the *Oracle Fusion Middleware Exalogic Machine Owner's Guide*.

On the firewall protecting the data tier, the database listener port (typically, `1521`) is required to be open. The LDAP ports (typically, `389` and `636`) are also required to be open for the traffic accessing the LDAP storage in the Exalogic enterprise deployment.

2.5 Load Balancer Requirements

This enterprise topology uses an external load balancer. This external load balancer should have the following features:

- Ability to load-balance traffic to a pool of physical servers through a virtual host name: Clients access services using the virtual host name (instead of using actual host names). The load balancer can then load balance requests to the servers in the pool.
- Port translation configuration should be possible, so that incoming requests on the virtual host name and port are directed to a different port on the backend servers.
- Monitoring of ports on the servers in the pool to determine availability of a service.
- Virtual servers and port configuration: Ability to configure virtual server names and ports on your external load balancer, and the virtual server names and ports must meet the following requirements:
  - The load balancer should allow configuration of multiple virtual servers. For each virtual server, the load balancer should allow configuration of traffic management on more than one port. For example, for Oracle HTTP Server in the web tier, the load balancer needs to be configured with a virtual server and ports for HTTP and HTTPS traffic.
  - The virtual server names must be associated with IP addresses and be part of your DNS. Clients must be able to access the external load balancer through the virtual server names.
- Ability to detect node failures and immediately stop routing traffic to the failed node.
- Fault-tolerant mode: It is highly recommended that you configure the load balancer to be in fault-tolerant mode.
- It is recommended that you configure the load balancer virtual server to return immediately to the calling client when the backend services to which it forwards traffic are unavailable. This is preferred over the client disconnecting on its own after a timeout based on the TCP/IP settings on the client machine.
- Sticky routing capability: Ability to maintain sticky connections to components. Examples of this include cookie-based persistence, IP-based persistence, and so on.
- The load balancer should be able to terminate SSL requests at the load balancer and forward traffic to the backend real servers using the equivalent non-SSL protocol (for example, HTTPS to HTTP).

Note: These are recommendations only. You can use your existing data center practices in the Oracle Exalogic environment.

2.6 Example: Horizontal Slicing Within Exalogic Machine Quarter Rack

Figure 2–3 illustrates horizontal slicing of an Oracle Exalogic machine quarter rack for Oracle Solaris and Oracle Linux physical environments.

Figure 2–3 Horizontal Slicing Within Exalogic Machine Quarter Rack

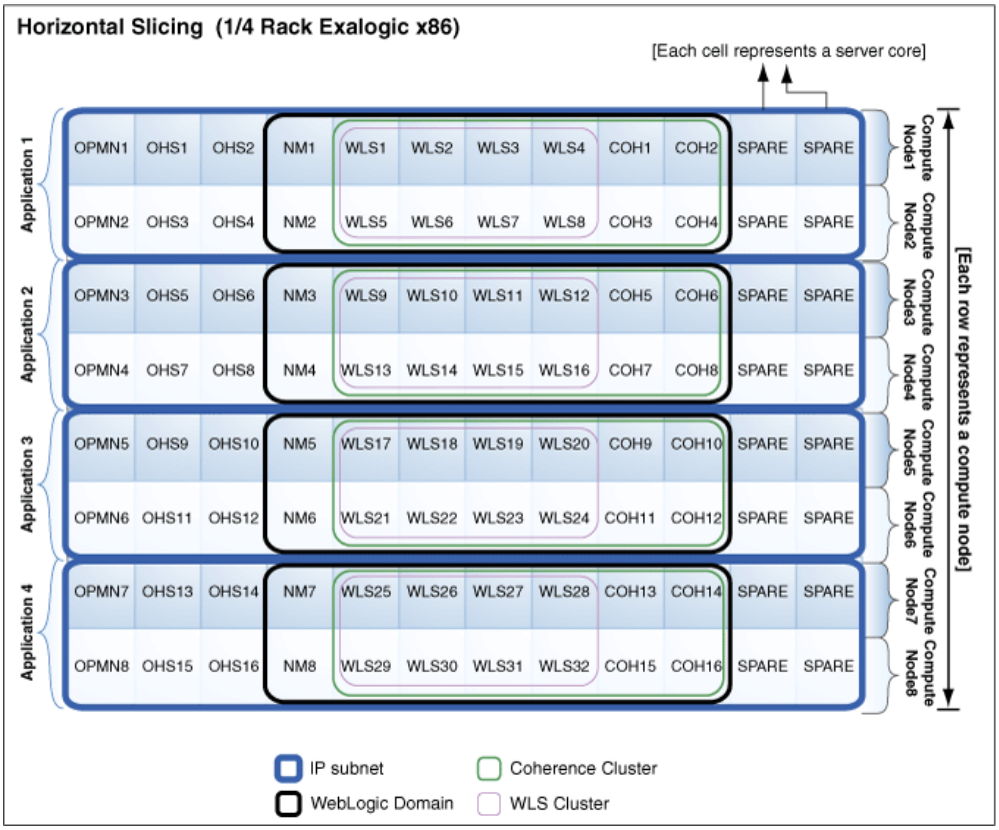

Based on this example, you can configure the following to deploy an application, such as `Dept_1`:

- A dedicated Node Manager instance per compute node
- A dedicated Oracle Process Manager and Notification Server (OPMN) process on each compute node
- Optional: two Oracle HTTP Server instances on each compute node if you wish to use Oracle HTTP Server inside Exalogic to load balance traffic on the IPoIB network
- Two storage-enabled Coherence Servers per compute node
- One WebLogic cluster with 8 Managed Servers across two compute nodes
- One Coherence cluster across two compute nodes

Note: The horizontal slicing of an Exalogic machine is an example only. All of the above are in the same IP subnet.

3

Network, Storage, and Database Preconfiguration

This chapter describes network, database, and storage preconfiguration required by the Oracle Exalogic enterprise deployment topology.

This chapter contains the following sections:

3.1 Important Notes for Oracle Solaris Users

If you are using the Oracle Solaris operating system on Exalogic compute nodes, keep the following points in mind:

- `BOND0` and `BOND1`, two important terms used in this guide, refer to the default interfaces for IP over InfiniBand (IPoIB) and Ethernet over InfiniBand (EoIB), respectively, on the Oracle Linux operating system.
- Oracle Solaris uses the IP Multipathing (IPMP) technology to support **IPMP Groups** that consist of one or more physical interfaces on the same system that are configured with the same IPMP group name. This technology provides the same functionality as *Bonded Interfaces* on Oracle Linux. You can name the IPMP groups anything. In this guide, `BOND0` and `BOND1` are used as example names to keep the terminology consistent with Oracle Linux.

IPMP Overview for Oracle Solaris Users

On the Oracle Solaris operating system, IP network multipathing (IPMP) provides physical interface failure detection and transparent network access failover for a system with multiple interfaces on the same IP link. IPMP also provides load spreading of packets for systems with multiple interfaces.

This section discusses the following topics:

IPMP Components

IPMP comprises the following components:

- The `in.mpathd` daemon
- The `/etc/default/mpathd` configuration file
- `ifconfig` options for IPMP configuration

> **Note:** For information about the `in.mpathd` daemon and the `mpathd` configuration file, see the *in.mpathd (1M)* man page on the Oracle Solaris operating system installed on Exalogic compute nodes. For information about `ifconfig`, see the *ifconfig (1M)* man page.

IPMP Groups

An IP multipathing group, or IPMP group, consists of one or more physical interfaces on the same system that are configured with the same IPMP group name. All interfaces in the IPMP group must be connected to the same IP link. The same (non-null) character string IPMP group name identifies all interfaces in the group. You can place interfaces from NICs of different speeds within the same IPMP group, as long as the NICs are of the same type. IPMP groups on Oracle Solaris provide the same functionality as Bonded Interfaces on Oracle Linux in the Exalogic environment. For example, the default IPMP group `ipmp0` comprises two physical interfaces that are connected to the default IPoIB link for internal communication in your Exalogic machine. The other default IPMP group `ipmp1` comprises two virtual interfaces that are connected to the default EoIB link for external data center connectivity.

> **Note:** For information about administering and configuring IPMP groups on the Oracle Solaris operating system installed on Exalogic compute nodes, see "Oracle Solaris 11 Express System Administrator Collection" at: `http://download.oracle.com/docs/cd/E19963-01/index.html`.

3.2 Machines

In this enterprise deployment guide, two compute nodes `ComputeNode1` and `ComputeNode2` are used as example compute nodes. They are located in Unit 2 and in Unit 3 of the Exalogic Machine rack, respectively.

One WebLogic Clusters (`Dept1_Cluster1`), containing 8 Managed Servers, will be configured to run on `ComputeNode1` and `ComputeNode2` in the example configuration scenario. Members of these WebLogic clusters are storage-disabled Coherence members.

> **Note:** However, in the configuration example described in this guide, the sample web application will be deployed to the Oracle WebLogic cluster (`Dept1_Cluster1`).

One Coherence clusters (`CoherenceCluster1`), containing Coherence servers and storage-disabled WLS servers, are created across `ComputeNode1` and `ComputeNode2`.

For information about Oracle Exalogic machine rack layout, see the topic "Exalogic Machine Rack Layout" in the *Oracle Fusion Middleware Exalogic Machine Owner's Guide*.

> **Note:** The default `NET0` IP addresses of these example compute nodes, assigned at the time of manufacturing, are `192.168.1.1` and `192.168.1.2`. These IP addresses are used as the example IP addresses of the physical machines `ComputeNode1` and `ComputeNode2` in this guide. These addresses will be used for access to server administration functions from the data center management Ethernet network.
>
> The default, factory-assigned InfiniBand `BOND0` (bonded interface for IP over InfiniBand connectivity) assigned to `ComputeNode1` and `ComputeNode2` (located in Unit 2 and Unit 3 of the Oracle Exalogic machine rack, respectively) are `192.168.10.1` and `192.168.10.2`, respectively. You may replace these IP addresses with your own IP addresses.

Table 3–1 Machine Names

Name Used in This Guide	Description
`LBR`	Represents external load balancer to distribute load across and failover between web servers.
`WEBHOST1`	Hosts a web server outside of the Oracle Exalogic machine.
`WEBHOST2`	Hosts a web server outside of the Oracle Exalogic machine.
`ComputeNode1`	Hosts Oracle WebLogic Server components, including Managed Servers and Administration Server. In addition, this machine hosts Coherence servers and Node Manager.
`ComputeNode2`	Hosts Oracle WebLogic Server components, including Managed Servers. In addition, this machine hosts Coherence servers and Node Manager.

> **Note:** The number of physical machines (compute nodes) in the Oracle Exalogic machine depends on your purchased hardware configuration. Oracle Exalogic machine full rack contains 30 compute nodes, an Oracle Exalogic machine half rack contains 16 compute nodes, and an Oracle Exalogic machine quarter rack contains 8 compute nodes.

3.3 Network

This section covers the following topics:

- Section 3.3.1, "General Network and InfiniBand Setup"
- Section 3.3.2, "Network Diagram for Exalogic Machine"
- Section 3.3.3, "Enterprise Deployment Network Configuration"
- Section 3.3.4, "Virtual Server Names"
- Section 3.3.5, "Load Balancers"
- Section 3.3.6, "Firewalls and Ports"

3.3.1 General Network and InfiniBand Setup

Before configuring the Exalogic enterprise deployment reference topology, be sure to complete the following prerequisites:

- Set up and configure the InfiniBand gateways and switches in Exalogic Machine, as described in the *Oracle Fusion Middleware Exalogic Machine Owner's Guide*.
- Run the Exalogic Configuration Utility to complete the initial network configuration, such as assignment of IP addresses, routing tables, and so on.
- Ensure that Subnet Manager (Master) requirements are satisfied, as described in "Subnet Manager Requirements for Connecting Exalogic Machine to Oracle Exadata Database Machine" section in the *Oracle Fusion Middleware Exalogic Machine Owner's Guide*.
- Ensure that the default IP over InfiniBand (`IPoIB`) and Ethernet over InfiniBand (`EoIB`) bonded interfaces are configured. The IPoIB bonded interface is configured by Exalogic Configuration Utility, by default. You must configure EoIB bonded interfaces manually.

Note: For more information about networking configuration, see the *Oracle Fusion Middleware Exalogic Machine Owner's Guide*.

For information about administrator roles and definitions, see Section 1.6, "Administrator Roles and Permissions".

3.3.2 Network Diagram for Exalogic Machine

Figure 3–1 shows the network diagram for an Oracle Exalogic machine.

Figure 3–1 Exalogic Machine Network Overview

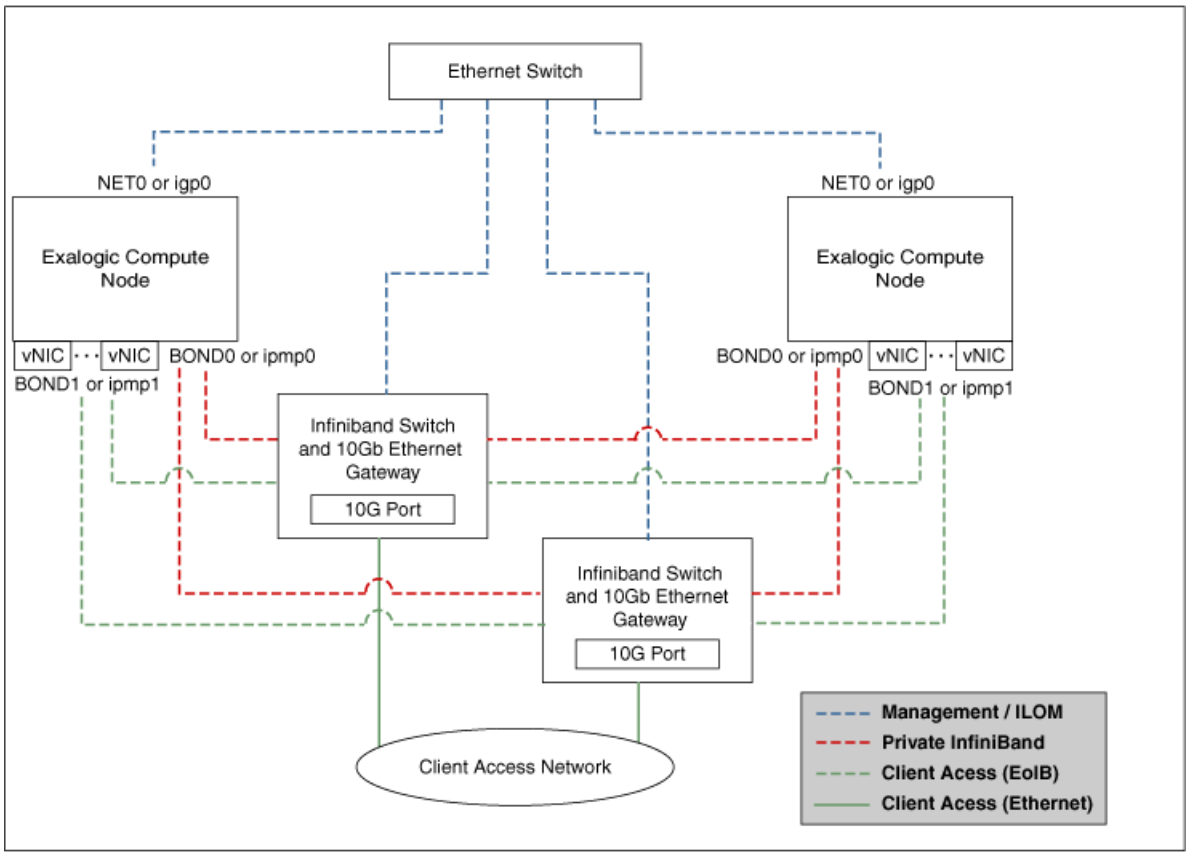

Diagram of the Oracle Exalogic machine network, which is described in detail in the text preceding and following the graphic.

****

The schematic representation of Oracle Exalogic machine's network connectivity includes the following:

- Default `BOND0` interface, which is the private InfiniBand fabric including the compute nodes connected via Sun Network QDR InfiniBand Gateway Switches

  Note: InfiniBand `BOND0` interfaces are the default channel of communication among Exalogic compute nodes and storage server head. IP subnets and additional bonds can be added on top of this default bonded interface.

  The device nodes representing the IPoIB network interface for Oracle Linux are referred to as `ib0` and `ib1`. The corresponding logical devices created by Oracle Solaris are referred to as `ibp0` and `ibp1`. The default IPoIB bonded interface `BOND0` or `IPMP0`, configured by the Exalogic Configuration Utility, comprises these Linux-specific interfaces or Solaris-specific interfaces, respectively.

- `BOND1` interface, which is the Ethernet over InfiniBand (EoIB) link

Note: The device nodes representing the EoIB network interface for Oracle Linux are referred to as `vnic0` and `vnic1`. The Linux kernel creates `eth` device nodes that correspond to the `vnic0` and `vnic1` instances that are created on the Sun Network QDR InfiniBand Gateway Switch.

The corresponding logical devices created by Oracle Solaris are referred to as `eoib0` and `eoib1`. The EoIB bonded interface `BOND1` or `IPMP1` must be configured manually. When you configure them, choose the network interfaces specific to your operating system.

- `NET0` interface, which is associated with the host Ethernet port 0 IP address for every compute node and storage server head

Note: The device node representing the management network interface for Oracle Linux is referred to as `eth0`. The corresponding logical device created by Oracle Solaris is referred to as `igb0`.

- Client access network for external data center connectivity

3.3.3 Enterprise Deployment Network Configuration

The WebLogic cluster (`Dept1_Cluster1`) and Coherence cluster require network configuration, such as creating a private subnet over the default IP over InfiniBand (IPoIB) interface. The subnet should support internal IP addresses, such as 10.x.x.x to the clusters for use by addresses that only need to be accessible to Oracle HTTP Server, load balancer, or other cluster members.

This section contains the following topics:

- Section 3.3.3.1, "IP Address and Network Channel Requirements"
- Section 3.3.3.2, "Determining Network Interface and Channel Requirements for a WebLogic Managed Server and the Administration Server"
- Section 3.3.3.3, "IP Addresses for Private InfiniBand Fabric Used by WebLogic Clusters and Coherence Clusters"
- Section 3.3.3.4, "IP Addresses for WebLogic Clusters When HTTP or T3 Traffic Is Via Ethernet over InfiniBand (EoIB)"
- Section 3.3.3.5, "Optional Network Configuration"

3.3.3.1 IP Address and Network Channel Requirements

The enterprise deployment configuration example (`Dept_1`) requires the following:

- **Private InfiniBand fabric network traffic used by WebLogic Server instances and Coherence servers** - The default channel of WebLogic Server, such as `7001`, listens to the private InfiniBand fabric network, which can multiplex multiple protocols, such as T3, HTTP, LDAP, and so on. This network and the WebLogic Server channel are meant for internal communication among WebLogic Server instances running on Exalogic compute nodes for communication like Managed Server-to-Administration Server traffic and cluster communication. This network supports both SDP and IP over InfiniBand. Coherence Servers typically also leverage this network.

For each WebLogic Managed Server, and the Administration Server, you should use a unique floating IP address(BOND0) for the default channel.

Note: In the enterprise deployment network configuration described in this guide, you are creating and assigning individual floating IP addresses (private) for Oracle WebLogic Managed Servers and the Administration Server. These floating IP addresses are configured using the default BOND0 interface and based on a suitable net mask. In the configuration example described in this guide, you are creating 16 Managed Servers, 8 Coherence Servers (For Coherence Servers, you are assigning the BOND0 IP address of the complete nodes on which the servers are running), and 1 WebLogic Administration Server for the Dept_1 application domain. Therefore, you should define a suitable range for the net mask to cover all of these servers in Dept_1.

This guide uses 10.0.0.0 as an example IP subnet for the InfiniBand fabric of Dept_1. A 5-bit net mask of 255.255.255.224 is used in this guide to address the floating IP address requirements for 17 servers (16 WebLogic Managed Servers, and 1 Administration Server. These IP addresses are used as the listen addresses for WebLogic Managed Servers and for the Administration Servers in this guide.

The default network channel uses IP over InfiniBand (IPoIB).

- **HTTP traffic coming from the external data center or the Internet to the Exalogic internal traffic** - HTTP traffic from the Internet or from the external data center comes in over 10 Gb Ethernet via one of the Sun Network QDR InfiniBand Gateway Switches. Then the traffic reaches the WebLogic Server instances running on an Exalogic compute node via the Ethernet over InfiniBand network of Exalogic Machine (EoIB).

  In this scenario, you must configure a separate set of floating IP addresses using the BOND1 interface for each of the WebLogic Managed Servers and for the Administration Server. In addition, you must create at least two additional network channels on the BOND1 interface (bonded interface comprising vnic0 and vnic1 on each compute node) for each of the WebLogic Managed Servers and for the Administration Server:

  – One channel for HTTP

  – One channel for T3

  For each WebLogic Managed Server, you can use the same floating IP and port combination (BOND1) for these channels. However, you must set the right protocols. For more information, see Section 5.12, "Configuring Network Channels for HTTP and T3 Clients via EoIB".

Note: In the enterprise deployment network configuration described in this guide, you are creating and assigning individual floating IP addresses for each of WebLogic Managed Servers and for the Administration Server. These floating IP addresses are configured using the default `BOND1` interface and based on a suitable net mask.

In the configuration example described in this guide, you are creating 16 Managed Servers and 1 WebLogic Administration Server for the `Dept_1` application domain. Therefore, you should define a suitable range for the net mask to cover all of these servers in `Dept_1`.

- **WebLogic Server replication traffic** - This traffic requires a custom replication channel that uses Socket Direct Protocol (SDP). You must set the outbound-enabled attribute to `true`, so all outbound replication traffic can use the replication channel.

  This channel uses SDP. SDP is an InfiniBand feature that can be used as an alternative to TCP/IP that reduces network latency and CPU utilization.

  For each WebLogic Managed Server, you can use the same `BOND0` floating IP for the replication channel. However, you must use different ports by specifying additional replication ports. For more information, see Section 5.12, "Configuring Network Channels for HTTP and T3 Clients via EoIB"

3.3.3.2 Determining Network Interface and Channel Requirements for a WebLogic Managed Server and the Administration Server

Table 3–2 helps you identify and determine which virtual interfaces and network channels are necessary for a WebLogic Managed Server and the Administration Server in your WebLogic administration domain.

This example is for `WLS1`, which is one of the Managed Servers in the `Dept1_Cluster1` cluster. You can use this information to determine network interface and channel requirements for other servers in the
Oracle WebLogic Clusters, as required.

Table 3–2 Summary of Interface and Channel Requirements for WLS1

Channel and Example	Interface	Protocol	Purpose	SDP Enabled?	Outbound Enabled?	Is Required in Exalogic?
Default channel	Floating IP for IPoIB `BOND0` FIP1 Port1 (floatingIP:port) For example, for `WLS1`, it is `10.0.0.1:7003`.	t3 and http	All network traffic not otherwise accounted for, on other channels. For example, traffic in `Dept1_Cluster1` among Managed Servers.	No	Not applicable	Yes
Replication channel (`ReplicationChannel1`)	Floating IP for IPoIB `BOND0` FIP1 Port2, Port3, Port4, and so on For example, for `WLS1`, it is `10.0.0.1` as the host and ports `7004`,`7005`,`7006`, and so on.	t3	Replication traffic	Yes	Yes	Only required when hosting a session-based web application with highly available sessions and the application is not using Coherence*Web
HTTP client channel (`HTTPClientChannel1`)	Floating IP for EoIB `BOND1` FIP1 Port 1 For example, for `WLS1`, it is `10.1.0.1:8001`.	http	Web application support	No	No	Only required if any HTTP clients use the 10 Gb Ethernet network for Exalogic incoming and outgoing traffic.
T3 client channel (`T3ClientChannel1`)	Floating IP for EoIB `BOND1` FIP1 Port 1 For example, for `WLS1`, it is `10.1.0.1:8003`.	t3	JMS/EJB/JMX/RMI client support	No	No	Only required if any T3/RMI clients use the 10 Gb Ethernet network for Exalogic incoming and outgoing traffic.

3.3.3.3 IP Addresses for Private InfiniBand Fabric Used by WebLogic Clusters and Coherence Clusters

The private InfiniBand fabric including WebLogic Clusters and Coherence clusters, presented in the enterprise deployment configuration example (`Dept_1`), requires a set of IP addresses for all WebLogic Managed Server-to-Administration Server traffic and for cluster communication. These IP addresses are associated with the `BOND0` interface.

Table 3–3 describes these IP address requirements. This table contains the following columns:

- **Name** - host name used in this enterprise deployment guide.
- **IP Name** - IP address name used in this guide.
- **Type** - type of IP address.
  - **Physical IP**: fixed to a single machine (compute node). In this guide, the `BOND0` IP addresses of `ComputeNode1` and `ComputeNode2` are referred to as the physical IP addresses.

  - **Floating IP**: assigned to a WebLogic Managed Server to allow for server migration.

    In the configuration example discussed in this guide, two floating addresses are used to each of the WebLogic Managed Servers and to the Administration Server. The first type of floating IP address uses the IPoIB (`BOND0`) interface, and the second type of floating IP uses the EoIB (`BOND1`) interface.

  - **Virtual IP**: used by a load balancer.

- **Host** - host where a corresponding IP address is used. For floating IP addresses, a range of hosts is given.
- **Bound By** - identifies which software component will use the corresponding IP address.
- **Scope** - shows where the corresponding IP address is resolved. *Cluster-scope* addresses only have to be resolvable by machines in the cluster. For a list of machines, see Section 3.2, "Machines". These addresses are only used for inter-cluster communication or for access by the load balancer. *Intranet-scope* addresses are used for internal purposes only.

Note: The configuration example for `Dept_1` uses an example IP subnet `10.0.0.0` for IPoIB. For assigning IP addresses to each of the WebLogic Managed Servers, Coherence Servers, Node Manager, and to the Administration Server, a 5-bit net mask `255.255.255.224` is used in the example. This IP address range provides approximately 30 IP addresses.

Table 3–3 Summary of IP Addresses (BOND0)

Name and Example Used in This Guide	IP Name	Type	Host	Bound By	Scope	Description
ADMINVHN1 Example IP address used in the guide: IPoIB (BOND0) - 10.0.0.17	FIP17	Floating	ComputeNode1	Administration Server	Cluster	A floating IP address for the Administration Server is recommended, if you want to manually migrate the Administration Server from ComputeNode1 to ComputeNode2.
ComputeNode1 Example IP address used in the guide: ComputeNode1 (BOND0) - 192.168.10.1	el01cn01-priv	Fixed	IP associated with the BOND0 interface for ComputeNode1	Node Manager	Cluster	BOND0 IP used by the Node Manager running on ComputeNode1. In the example described in this guide, a compute node in Unit 2 of the Oracle Exalogic machine rack is referred to as ComputeNode1.

Table 3–3 (Cont.) Summary of IP Addresses (BOND0)

Name and Example Used in This Guide	IP Name	Type	Host	Bound By	Scope	Description
ComputeNode2 Example IP address used in the guide: ComputeNode2 (BOND0) - 192.168.10.2	el01cn02-priv	Fixed	IP associated with the BOND0 interface for ComputeNode2	Node Manager	Cluster	BOND0 IP used by the Node Manager running on ComputeNode2. In the example described in this guide, a compute node in Unit 3 of the Oracle Exalogic machine rack is referred to as ComputeNode2.

Table 3–3 (Cont.) Summary of IP Addresses (BOND0)

Name and Example Used in This Guide	IP Name	Type	Host	Bound By	Scope	Description
BOND0 (IPoIB): WLS1(10.0.0.1) WLS2 (10.0.0.2) WLS3 (10.0.0.3) WLS4 (10.0.0.4) WLS5 (10.0.0.5) WLS6 (10.0.0.6) WLS7 (10.0.0.7) WLS8 (10.0.0.8) WLS9 (10.0.0.9) WLS10 (10.0.0.10) WLS11 (10.0.0.11) WLS12 (10.0.0.12) WLS13 (10.0.0.13) WLS14 (10.0.0.14) WLS15 (10.0.0.15) and WLS16 (10.0.0.16)	FIP1, FIP2, FIP3, FIP4, FIP5, FIP6, FIP7, FIP8, FIP9, FIP10, FIP11, FIP12, FIP13, FIP14, FIP15, and FIP16, respectively	Floating	ComputeNode1 and ComputeNode2	WebLogic Managed Servers running on ComputeNode1 and ComputeNode2	Cluster	All of the Managed Servers require server migration between ComputeNode1 and ComputeNode2.

Table 3–3 (Cont.) Summary of IP Addresses (BOND0)

Name and Example Used in This Guide	IP Name	Type	Host	Bound By	Scope	Description
BOND0 (IPoIB) of the compute nodes on which Coherence Servers are running: Coh1 (192.168.10.1) Coh2 (192.168.10.1) Coh3 (192.168.10.1) Coh4 (192.168.10.1) Coh5 (192.168.10.2) Coh6 (192.168.10.2) Coh7 (192.168.10.2) and Coh8 (192.168.10.2)	BOND0 IP of ComputeNode1 BOND0 IP of ComputeNode2, respectively	Physical	ComputeNode1 and ComputeNode2	Coherence servers (nodes) spanning WebLogic cluster (Dept1_Cluster1)	Cluster	Coherence servers do not require migration between ComputeNode1 and ComputeNode2. In this guide, example IP BOND0 addresses of compute nodes are used.
WEBHOST1	IP3	Fixed	WEBHOST1	OHS1	Cluster	Oracle HTTP Server (OHS) is external to Exalogic Machine.
WEBHOST2	IP4	Fixed	WEBHOST2	OHS2	Cluster	
exalogic.mycompany.com	VIP26	Virtual	Load Balancer	Public	Public	External access point to WebLogic cluster (Dept1_Cluster1).
admin.mycompany.com	VIP27	Virtual	Load Balancer	Load Balancer	Intranet	Internal access to WebLogic Server Administration Console.
exalogicinternal.mycompany.com	VIP28	Virtual	Load Balancer	Load Balancer	Intranet	Internal access to WebLogic cluster (Dept1_Cluster1).

Floating IP addresses are IP addresses that may be re-assigned between compute nodes in the cluster. For example, if ComputeNode1 fails or goes down, then WebLogic Managed Servers running on ComputeNode1 (WLS1, WLS2, WLS3, WLS4, WLS5, WLS6, WLS7, and WLS8) can be migrated to ComputeNode2. In this case, you must activate the floating IPs of these Managed Servers on the new target machine ComputeNode2. When Node Manager is set up, it manages the registration and

removal of the floating IP addresses with the exception of the Administration Server's floating IP address.

Note that Managed Servers require separate floating IP addresses. You do not need to use the same host names used in this guide. The host names must be distinct. You cannot use a single IP address as both fixed IP and floating IP.

Note: Cluster-scope IP addresses must be in the `/etc/hosts` file of all Exalogic compute nodes. They are private to the InfiniBand fabric involving WebLogic clusters and Coherence clusters. Intranet-scope IP addresses must be available on the internal DNS servers. Public-scope IP addresses must be available on both external and internal DNS servers.

3.3.3.4 IP Addresses for WebLogic Clusters When HTTP or T3 Traffic Is Via Ethernet over InfiniBand (EoIB)

If any HTTP, T3, or RMI clients use the 10 Gb Ethernet network for Exalogic, you should configure virtual interfaces for WebLogic Managed Servers and for the Administration Server using the `BOND1` interface (`EoIB`).

Table 3–4 describes these IP address requirements.

Note: The configuration example for `Dept_1` uses an example IP subnet `10.1.0.0` for EoIB. For assigning IP addresses to each of the WebLogic Managed Servers and to the Administration Server, a 5-bit net mask `255.255.255.224` is used in the example. This IP address range provides approximately 30 IP addresses.

Table 3–4 Summary of IP Addresses (BOND1)

Name and Example Used in This Guide	IP Name	Type	Host	Bound By
ADMINVHN1 Example IP address used in the guide: EoIB (BOND1) - 10.1.0.17	FIP17	Floating	ComputeNode1	Administration Server A floating IP address for the Administration Server is recommended, if you want to migrate the Administration Server manually from ComputeNode1 to ComputeNode2.
BOND1 (EoIB): WLS1(10.1.0.1) WLS2 (10.1.0.2) WLS3 (10.1.0.3) WLS4 (10.1.0.4) WLS5 (10.1.0.5) WLS6 (10.1.0.6) WLS7 (10.1.0.7) WLS8 (10.1.0.8) WLS9 (10.1.0.9) WLS10 (10.1.0.10) WLS11 (10.1.0.11) WLS12 (10.1.0.12) WLS13 (10.1.0.13) WLS14 (10.1.0.14) WLS15 (10.1.0.15) and WLS16 (10.1.0.16)	FIP1, FIP2, FIP3, FIP4, FIP5, FIP6, FIP7, FIP8, FIP9, FIP10, FIP11, FIP12, FIP13, FIP14, FIP15, and FIP16, respectively	Floating	ComputeNode1 and ComputeNode2	WebLogic Managed Servers running on ComputeNode1 and ComputeNode2

3.3.3.5 Optional Network Configuration

This section also discusses the following topics:

- Section 3.3.3.5.1, "Application Isolation by IP Subnetting over IPoIB"
- Section 3.3.3.5.2, "Allowing a Compute Node to Access Two Different Subnets Simultaneously"

3.3.3.5.1 Application Isolation by IP Subnetting over IPoIB The configuration example described in this document uses two compute nodes ComputeNode1 and ComputeNode2 which are used by Dept_1 for deploying a web application, such as dizzyworld.ear.

To create an IP subnet over this default IPoIB link (BOND0) to isolate Dept_1 from other departments using the remaining Exalogic compute nodes, do the following:

1. Ensure that all of the compute nodes are configured with valid IP addresses.
2. Determine an IP address range for the subnet you are trying to create. In this example, an IP subnet 10.0.0.0 is used. The configuration example for Dept_1 requires at least 17 usable IP addresses using the BOND0 interface (16 Managed Servers and 1 Administration Server). Pick a range that provides you with approximately 30 addresses.

3. Calculate a net mask for this IP address range. For example, the sample net mask for a subnet to cover the above example IP addresses is `255.255.255.224`.
4. Log in to `ComputeNode1` as `root`.
5. Run `/usr/sbin/setup` file.

   The setup screen is displayed.
6. Select the `Network Configuration` option.
7. Use Tab key and select **Run Tool**, and then enter **Return**.
8. Use up and down arrows to select **Edit Devices**, and then hit return.
9. Use up and down arrows to select **bond0**, and then hit return.
10. Type the respective IP address and net mask. For example, enter `255.255.255.224` as the net mask. The example IP subnet used is `10.0.0.0`.
11. Restart the appropriate network interfaces using the `ifconfig` command as follows:

```
# ifconfig bond0 <IP_address> <net_mask> up
```

> **Note:** Similarly, you can create IP subnet `10.0.0.32` for `Dept_2`, IP subnet `10.0.0.64` for `Dept_3`, and so on.

3.3.3.5.2 Allowing a Compute Node to Access Two Different Subnets Simultaneously If you isolated the application deployment and environment of the `Dept_1` department from another department, such as `Dept_2`, then you must create separate IP subnets for both `Dept_1` and `Dept_2` over the default IP over InfiniBand (IPoIB) link. `Dept_1` uses `ComputeNode1` and `ComputeNode2`. `Dept_2` uses `ComputeNode3` and `ComputeNode4`.

In some scenarios, the `Dept_1` application may require communication with the `Dept_2` application. To enable the `Dept_1` application (deployed on `ComputeNode1` and `ComputeNode2`) to communicate with the `Dept_2` application (deployed on `ComputeNode3` and `ComputeNode4`), you must set up another IP subnet in which both compute nodes of `Dept_1` and `Dept_2` are members.

This example uses two IP subnets: `10.0.0.0` for `Dept_1`, and `10.0.0.32` for `Dept_2`

You should create child network interfaces (referred to as **logical interfaces** on Oracle Solaris) for the `Dept_1` compute nodes in a subnet that require connectivity to `Dept_2` compute nodes, which are in a different subnet. Child interfaces or logical interfaces are created using standard IP aliasing. Oracle recommends that you identify such connection and access requirements for the department-level applications deployed on Exalogic compute nodes when configuring the network for enterprise deployment. In addition, you should add IP aliasing in your network configuration files to avoid reconfiguration.

Example Scenario

Consider an IPoIB network with four compute nodes: `ComputeNode1`, `ComputeNode2`, `ComputeNode3`, and `ComputeNode4`.

`ComputeNode1` and `ComputeNode2` are part of a cluster defined by IP subnet `10.0.0.0`, while `ComputeNode3` and `ComputeNode4` are part of a cluster defined by IP subnet `10.0.0.32`. Both subnets use `255.255.255.224` as their net masks. In

this example, the WLS1 Managed Server with a floating IP 10.0.0.1 is running on a compute node in the first subnet (10.0.0.0).

- Run the following commands after logging in to ComputeNode1 or ComputeNode2 as a root user:

  ```
  # ifconfig bond0:x 10.0.0.x netmask 255.255.255.224 up
  ```

  where 10.0.0.x is either 10.0.0.1 or 10.0.0.2, the IPs of ComputeNode1 and ComputeNode2, respectively.

- Configure two additional nodes: ComputeNode3 and ComputeNode4

  Run the following commands after logging in to ComputeNode3 or ComputeNode4 as a root user:

  ```
  # ifconfig bond0:x 10.0.0.z netmask 255.255.255.224 up
  ```

  where 10.0.0.z is either 10.0.0.33 or 10.0.0.34, the IPs of ComputeNode3 and ComputeNode4, respectively.

The four nodes are now on two separate IP subnets. ComputeNode1 and ComputeNode2 are on subnet 10.0.0.0, while ComputeNode3 and ComputeNode4 are on subnet 10.0.0.32.

Now, a Managed Server in the first subnet requires access to a compute node on the second subnet. To achieve this, configure an additional subnet 10.0.10.0 with net mask 255.255.255.192. In addition, perform the following configuration:

```
# ifconfig bond0:1 10.0.10.y netmask 255.255.255.192 up
```

where 10.0.10.y is either 10.0.10.1 or 10.0.10.2, the new IPs of ComputeNode1 and ComputeNode2 in the new subnet, respectively.

> **Note:** On Oracle Solaris, you must first plumb the logical interface by running the command `ifconfig bond0:1 plumb`. Then you must run the `ifconfig bond0:1 10.0.10.y netmask 255.255.255.192 up` command to start the logical interface.

Run the following command:

```
# ifconfig bond0:1 10.0.10.w netmask 255.255.255.192
```

where 10.0.10.w is either 10.0.10.3 or 10.0.10.4, the IPs of ComputeNode3 and ComputeNode4 in the new subnet, respectively.

These commands add the network routing entries that are required to enable a compute node to access two different subnets simultaneously.

3.3.4 Virtual Server Names

The Exalogic enterprise deployment reference topology uses the following virtual server names:

- Section 3.3.4.1, "exalogic.mycompany.com"
- Section 3.3.4.2, "admin.mycompany.com"
- Section 3.3.4.3, "exalogicinternal.mycompany.com"

Ensure that the virtual server names are associated with IP addresses and are part of your DNS configuration. The nodes running Oracle Fusion Middleware must be able to resolve these virtual server names.

Note: This document does not discuss DNS configuration.

3.3.4.1 exalogic.mycompany.com

`exalogic.mycompany.com` is a virtual server name that acts as the external access point to the Oracle WebLogic cluster (`Dept1_Cluster1`). This virtual server is defined on the load balancer, and it maps to a VIP (load balancer). For more information, see Table 3.3.3, "Enterprise Deployment Network Configuration".

3.3.4.2 admin.mycompany.com

`admin.mycompany.com` is an example virtual server name that provides internal access to the Oracle WebLogic Server Administration Console. You can define this virtual server on the management LAN load balancer or direct the traffic to a web server instance meant for handling internal traffic. This virtual server name uses a VIP assigned to the load balancer to fail over to another compute node in Exalogic Machine. For more information, see Table 3.3.3, "Enterprise Deployment Network Configuration".

3.3.4.3 exalogicinternal.mycompany.com

`exalogicinternal.mycompany.com` is a virtual server name that acts as the internal access point to the Oracle WebLogic cluster (`Dept1_Cluster1`). This virtual server is defined on the load balancer, and it uses a VIP assigned to the load balancer to distribute load across several compute nodes in Exalogic Machine. For more information, see Table 3.3.3, "Enterprise Deployment Network Configuration".

3.3.5 Load Balancers

This enterprise topology uses an external load balancer. For more information about load balancers, see Section 2.4.1, "Web Tier."

Note: The Oracle Technology Network (`http://www.oracle.com/technology/index.html`) provides a list of validated load balancers and their configuration at `http://www.oracle.com/technetwork/middleware/ias/tested-lbr-fw-sslaccel-100648.html`.

The number of Oracle HTTP Server instances depends on the WebLogic Server instances you require for deploying your application. This example configuration uses two Oracle HTTP Server instances `WEBHOST1` and `WEBHOST2` only.

3.3.5.1 Configuring the Load Balancer

To configure the load balancer, complete these steps:

1. Create two pools of servers. You will assign these pools to virtual servers.
2. Add the addresses of the first set of Oracle HTTP Server (https) hosts to one pool. For example:
   - `WEBHOST1`:4443
   - `WEBHOST2`:4443
3. Add the addresses of the second set of Oracle HTTP Server (http) hosts to another pool. For example:

- WEBHOST1:7777
- WEBHOST2:7777

4. Configure a virtual server in the load balancer for exalogic.mycompany.com:443.
   - For this virtual server, use your Oracle Exalogic machine's frontend address as the virtual server address (for example, exalogic.mycompany.com). The frontend address is the externally facing host name used by your Oracle Exalogic machine and that will be exposed in the Internet.
   - Configure this virtual server with port 80 and port 443. Any request that goes to port 80 should be redirected to port 443.
   - Enable address and port translation.
   - Enable reset of connections when services and/or nodes are down.
   - Assign the pool, created in step 2 or 3, to the virtual server,. This depends on whether SSL terminates at the load balancer or passes through.
   - Create rules to filter out access to /console on this virtual server. For more information, see the load balancer documentation provided by your vendor.
5. Configure a virtual server in the load balancer for admin.mycompany.com:80.
   - For this virtual server, use your Oracle Exalogic machine's internal administration address as the virtual server address (for example, admin.mycompany.com). This address is typically not externalized.
   - Specify HTTP as the protocol.
   - Enable address and port translation.
   - Enable reset of connections when services and/or nodes are down.
   - Assign the pool created in step 1 to the virtual server.
6. Configure monitors for the Oracle HTTP Server nodes to detect failures in these nodes.
   - Set up a monitor to regularly ping the "/" URL context.

     Tip: Use GET /\n\n instead if the Oracle HTTP Server's document root does not include index.htm and Oracle WebLogic Server returns a 404 error for "/".

   - For the ping interval, specify a value that does not overload your system. You can try 5 seconds as a starting point.
   - For the timeout period, specify a value that can account for the longest time response that you can expect from your Oracle Fusion Middleware product application, that is, specify a value greater than the longest period of time any of your requests to HTTP servers can take.

3.3.6 Firewalls and Ports

Many Oracle Fusion Middleware components and services use ports. As an administrator, you must know the port numbers used by these services, and to ensure that the same port number is not used by two services on a host.

Most port numbers are assigned during installation.

Table 3–5 lists the ports used in the Oracle Exalogic deployment reference topology, including the ports that you must open on the firewalls in the topology.

Firewall notation:

- FW0 refers to the outermost firewall.
- FW1 refers to the firewall between the web tier and the application tier.
- FW2 refers to the firewall between the application tier and the data tier.

Note: If you are connecting an Oracle Exalogic machine to Oracle Exadata Database Machine to run on the same InfiniBand fabrics for database connectivity, this firewall (FW2) does not apply.

For more information about port numbers, see the topic "Port Numbers by Component" in the *Oracle Fusion Middleware Administrator's Guide*.

Table 3–5 Ports Used in the Reference Topology

Type	Firewall	Port and Port Range	Protocol / Application	Inbound / Outbound	Other Considerations and Timeout Guidelines
Browser request	FW0	80	HTTP / Load Balancer	Inbound	Timeout depends on all HTML content and the process models used for the Oracle Fusion Middleware products you are using in the Exalogic environment.
Browser request	FW0	443	HTTPS / Load Balancer	Inbound	Timeout depends on all HTML content and the process models used for the Oracle Fusion Middleware products you are using in the Exalogic environment.
Load balancer to Oracle HTTP Server	n/a	7777 as the example HTTP port for `WEBHOST1` and `WEBHOST2`. 4443 as the example HTTPS port for `WEBHOST1` and `WEBHOST2`.	HTTP/HTTPS	n/a	See Section 3.3.5.1, "Configuring the Load Balancer." For actual values, see the topic "Port Numbers by Component" in the *Oracle Fusion Middleware Administrator's Guide*.

Table 3–5 (Cont.) Ports Used in the Reference Topology

Type	Firewall	Port and Port Range	Protocol / Application	Inbound / Outbound	Other Considerations and Timeout Guidelines
Administration Console access	FW1	7001	HTTP / Administration Server and Enterprise Manager	Both	You should tune this timeout based on the type of access to the admin console (whether it is planned to use the Oracle WebLogic Server Administration Console from application tier clients or clients external to the application tier).
Coherence	n/a	8088 Range: 8080 - 8090		n/a	n/a
Application tier to data tier (Oracle database or RAC outside of Oracle Exalogic machine via Ethernet)	FW2	1521		n/a	n/a
Managed Server Access (WLS1, WLS2, WLS3. WLS4, WLS5, WLS6, WLS7, WLS8, WLS9, WLS10, WLS11, WLS12, WLS13, WLS14, WLS15, and WLS16)	FW1	8001	HTTP	Inbound	Managed Servers, which use `BOND1` floating IP addresses, are accessed via Oracle HTTP Server.

3.4 Shared Storage and Recommended Project and Share Structure

The following section details the project and share structure that Oracle recommends for the Oracle Exalogic enterprise deployment topology. Other directory layouts are possible and supported, but the model adopted in this guide is chosen for maximum availability, providing both the best isolation of components and symmetry in the configuration. The rest of the document uses this directory structure and directory terminology.

File systems, including product binaries and Oracle Home directories for domain and products, are mounted via NFS from Sun ZFS Storage 7320 appliance, which is included in all Oracle Exalogic machine configurations.

This section covers these topics:

- Section 3.4.1, "Overview of Storage Configuration"
- Section 3.4.2, "Setting Up Enterprise Deployment Storage Configuration"
- Section 3.4.3, "Recommendation About Storage Location for Syslogs and Operating System Patches"

3.4.1 Overview of Storage Configuration

Storage is configured in **Pools** that are characterized by their underlying data redundancy, and provide space that is shared across all file systems and LUNs. During the configuration process, you will select which devices to allocate to a storage pool

and the redundancy profile most appropriate to your workload, balancing performance, availability, and capacity.

Note: For more information, see the Sun ZFS Storage documentation at the following URL:

http://download.oracle.com/docs/cd/E22471_01/index.html

In the Oracle Exalogic environment, the enterprise deployment reference topology uses the following:

- Storage disks on the Sun ZFS Storage 7320 appliance are allocated to a single storage pool, such as `exalogic`, by default. Every compute node in an Oracle Exalogic machine can access both of the server heads of the storage appliance. The storage pool uses one of the server heads, which are also referred to as controllers. The server heads use active-passive cluster configuration. Exalogic compute nodes access the host name or IP address of a server head and the mount point based on the distribution of the storage pool.
- If the active server head fails, the passive server head imports the storage pool and starts to offer services. From the compute nodes, the user may experience a pause in service until the storage pool is started to be serviced from the working server head. However, this delay may not affect client activity; it may affect disk I/O only.
- By default, data is mirrored, which yields a highly reliable and high-performing system. The default storage configuration is done at the time of manufacturing, and it includes the following shares:
  - Two exclusive NFS shares for each of the Exalogic compute nodes - one for crash dumps, and another for general purposes.

    In this scenario, you can implement access control for these shares, based on your requirements.
  - Two common NFS shares to be accessed by all compute nodes - one for patches, and another for general purposes.
- In addition, you can create department-level file systems referred to as **Projects**, such as `Dept_1`.
- Each department gets 2 compute nodes. For example, `Dept_1` in the enterprise deployment reference topology uses compute nodes `ComputeNode1` and `ComputeNode2`. All compute nodes can access the common NFS shares.
- Each department-level Project can be further divided into multiple partitions called **Shares**. For example, for `Dept_1`, you can create one Share for the Domain Home (`domains`), one optional share for JMS persistence logs and JTA logs (`jmsjta`).
- For file systems shared across compute nodes in Exalogic Machine, you can create individual projects. For example, the `FMW_Product1` project includes shares for Oracle WebLogic product binaries, logs, and configurations. If necessary, you can create and maintain separate file systems or shares for multiple WebLogic installations. You can use this method if you wish to maintain different versions of Oracle WebLogic Server (at different patch levels), based on your specific installation and management requirements.

This section also contains the following topics:

3.4.1.1 Viewing Default Storage Pool, Projects, and Shares

To view the default storage pool, projects, and shares created on the Sun ZFS Storage 7320 appliance in your Oracle Exalogic Machine, complete the following steps:

1. Ensure that IP address, host name, and network for the Sun ZFS Storage 7320 appliance are configured correctly.

   Note: This initial configuration is described in the "Initial Configuration of Oracle Exalogic machine Using Exalogic Configuration Utility" chapter and in the "Configuring the Sun ZFS Storage 7320 appliance" chapter in the *Oracle Fusion Middleware Exalogic Machine Owner's Guide*.

2. In the address bar of a browser, enter the IP address or host name for the storage appliance, which you assigned to the `NET0` port during the initial configuration of your Exalogic Machine as follows:

   ```
   https://ipaddress:215
   ```

   or

   ```
   https://hostname:215
   ```

   Note: For example, if you retain the default IP addresses assigned to Exalogic Machine components, you can access the server heads of the storage appliance by using either `192.168.1.115` or `192.168.1.116`.

   The login screen appears.

3. Type `root` in the **Username** field and the administrative password that you entered into the appliance shell kit interface and press the `Enter` key. The Welcome screen appears.

4. On the home page, click **Configuration**. Then click **STORAGE**. The default pool configuration is shown, as in Figure 3–2.

Figure 3–2 Default Pool Configuration

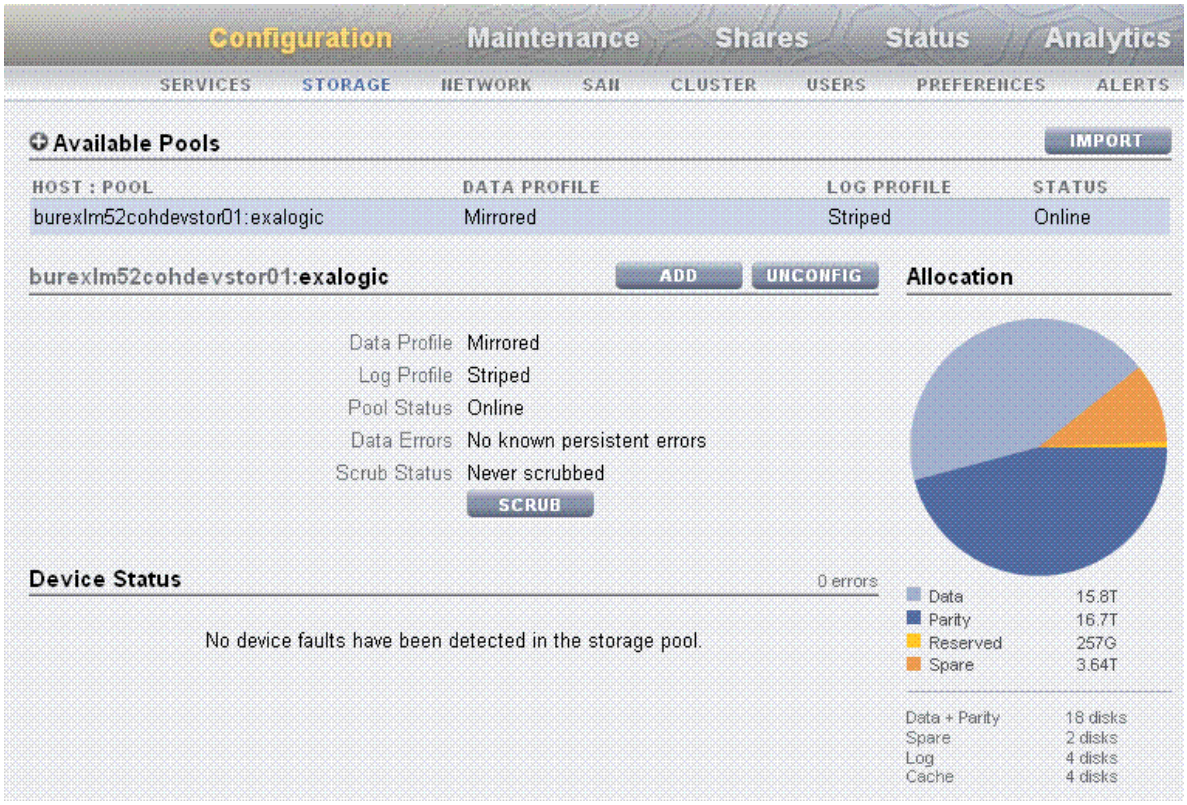

5. To view the default projects (common projects that can be accessed by all of the compute nodes and exclusive projects for each compute node), click **Shares** in Figure 3–2. Then click **PROJECTS**. The default projects and time of their creation are shown, as in Figure 3–3.

Figure 3–3 Default Projects

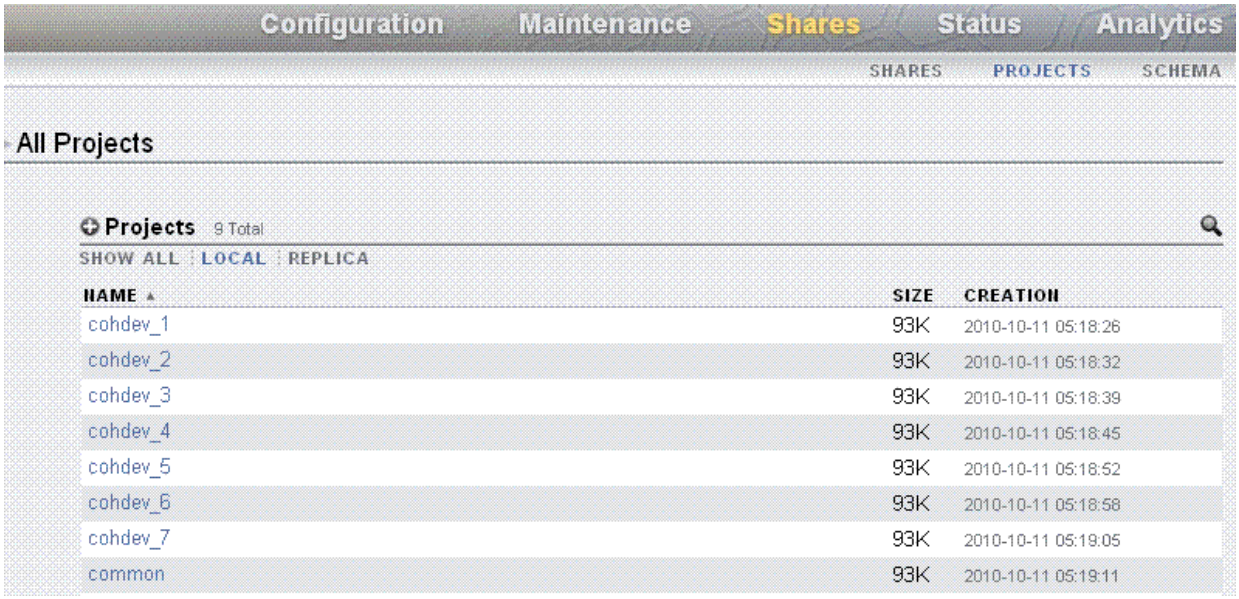

6. To view the default shares (common shares that can be accessed by all of the compute nodes and exclusive shares for each compute node), click **Shares** in Figure 3–2. Then click **SHARES**. The default shares and storage mount points are shown, as in Figure 3–4.

Figure 3–4 Default Shares

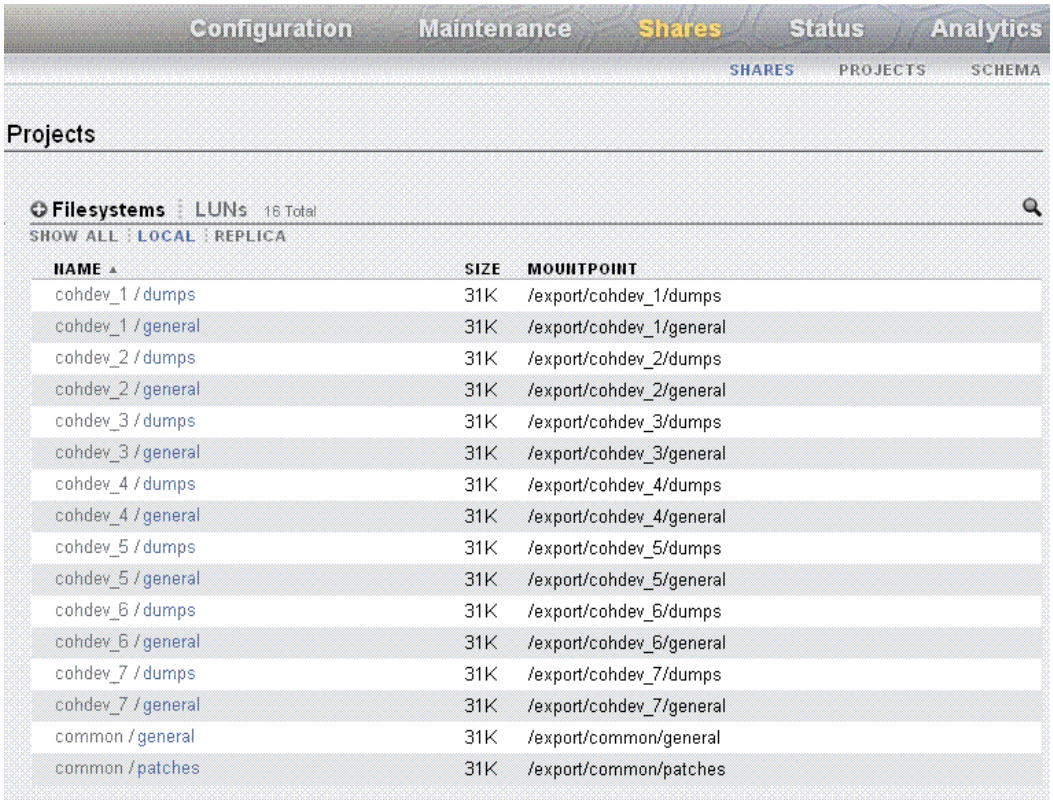

3.4.1.2 Recommended Project and Share Structure

Figure 3–5 illustrates the recommended project and share structure.

Figure 3–5 Project and Share Structure

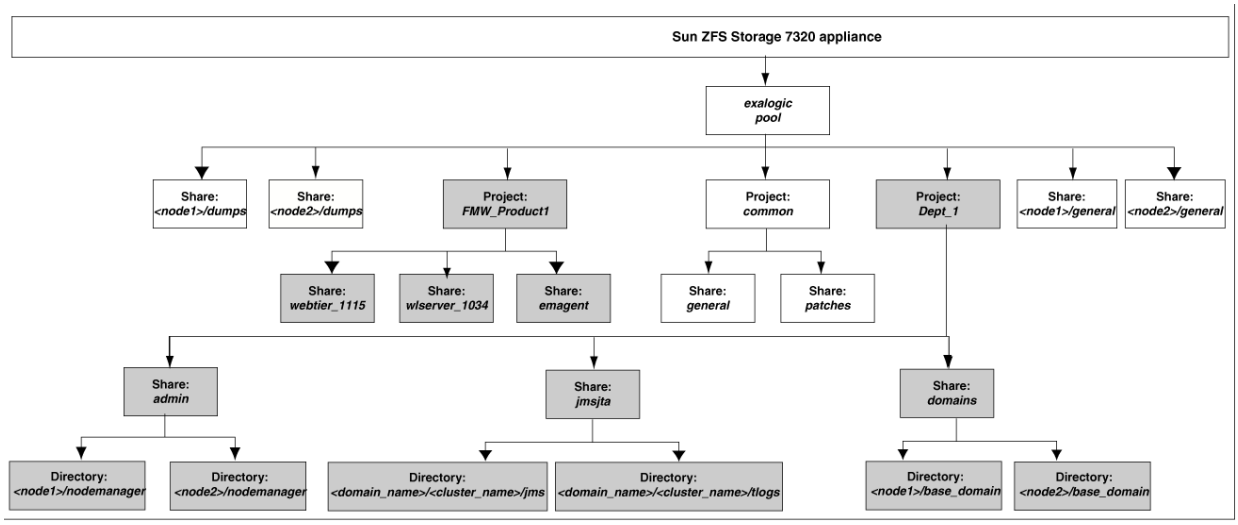

Note the following:

- Boxes without fill in the above diagram represent the default projects and shares that are created at the time of manufacturing. The project named `common` can be accessed by all of the compute nodes in the Oracle Exalogic machine.

  Note: The default shares `<node1>/dumps`, `<node1>/general`, `<node2>/dumps`, and `<node2>/general` are specific to `ComputeNode1` and `ComputeNode2`, which are the physical machines used by `Dept_1`.

- Boxes filled with gray represent custom projects and shares that you create based on your deployment requirements. The configuration example uses `Dept_1`. Therefore, the `Dept_1` project is created, and it can be accessed by `ComputeNode1` and `ComputeNode2`. Custom shares can include product binaries, logs, and domain home directories.

  Note: Figure 3–5 does not show other required internal directories, such as `jrockit`.

  In this document, the configuration example shows how to set up shared storage for one department, `Dept_1`, including two compute nodes `ComputeNode1` and `ComputeNode2`. The names of Managed Servers, physical IP addresses of compute nodes are used accordingly. You can extrapolate the example for creating a directory structure for your specific requirements related for application deployment, for security restrictions at file system level, and for specific SLAs.

- The `/u01/app/FMW_Product1/Oracle` directory is referred to as `ORACLE_BASE`.
- The `/u01/app/FMW_Product1/Oracle/Middleware` directory is referred to as `MIDDLEWARE_HOME`.

- The emagent share under the FMW_Product1 project is required only if you are using Oracle Enterprise Manager Grid Control 11*g* to monitor the topology.
- The webtier_1115 share under the FMW_Product1 is required only if you want to use Oracle HTTP Server to load balance traffic on Exalogic's BOND0/IPMP0 network.

3.4.2 Setting Up Enterprise Deployment Storage Configuration

Setting up enterprise deployment storage configuration for the Dept_1 example described in this guide involves the following steps:

- Creating Projects for Dept_1 and FMW_Product1
- Creating Shares for Dept_1
- Creating Shares for FMW_Product1
- NFSv4 Configuration Requirements
- Creating Mount Points on ComputeNode1 and ComputeNode2
- Editing the /etc/fstab (Linux) or /etc/vfstab (Solaris) File
- Mounting the Volumes
- Creating Groups and Users and Controlling Access to Mounted Shares

3.4.2.1 Creating Projects for Dept_1 and FMW_Product1

In the configuration example, you are creating a project for Dept_1 and a separate project for FMW_Product1. Dept_1 will contain department-level and machine-level shares, and FMW_Product1 will contain separate shares for storing product binaries, such as wlserver_10.3.

In the Browser User Interface (BUI), you can access the Projects user interface by clicking **Configuration** > **STORAGE** > **Shares** > **Projects**. The Project Panel is displayed.

To create the Dept_1 project, do the following:

1. Click the **+** button above the list of projects in the Project Panel.
2. Enter a name for the project, such as Dept_1. The new project Dept_1 is listed on the Project Panel, which is on the left navigation pane.
3. Click the **General** tab on the Dept_1 project page to set project properties. This section of the BUI controls overall settings for the project that are independent of any particular protocol and are not related to access control or snapshots. While the CLI groups all properties in a single list, this section describes the behavior of the properties in both contexts.

   The project settings page contains three sections: Space Usage (Users and Groups), Inherited Properties, and Default Settings (Filesystems and LUNs). Table 3–6 describes the project settings.

Table 3–6 Project Settings

Section and Setting	Description
Space Usage	Space within a storage pool is shared between all shares. Filesystems can grow or shrink dynamically as needed, though it is also possible to enforce space restrictions on a per-share basis. ■ **Quota** - Sets a maximum limit on the total amount of space consumed by all filesystems and LUNs within the project. ■ **Reservation** - Guarantees a minimum amount of space for use across all filesystems and LUNs within the project.
Inherited Properties	Standard properties that can either be inherited by shares within the project. The behavior of these properties is identical to that at the shares level. ■ **Mountpoint** - The location where the filesystem is mounted. This property is only valid for filesystems. Oracle recommends that you use specify `/export/<project_name>` as the default mountpoint. By using this consistently, you can group all shares and mount under the relevant project. It also prevents multiple shares from using the same mount points. Note that the same storage appliance is used by a multiple departments (15 in the case of Oracle Exalogic machine full rack configuration). The departments will have a similar share structure, such as `/export/dept_1/<share1>`, `/export/dept_2/share1`, and so on. ■ **Readonly** - Controls whether the filesystem contents are read only. This property is only valid for filesystems. ■ **Update access time on read** - Controls whether the access time for files is updated on read. This property is only valid for filesystems. ■ **Non-blocking mandatory locking** - Controls whether CIFS locking semantics are enforced over POSIX semantics. This property is only valid for filesystems. ■ **Data deduplication** - Controls whether duplicate copies of data are eliminated. ■ **Data compression** - Controls whether data is compressed before being written to disk. ■ **Checksum** - Controls the checksum used for data blocks. ■ **Cache device usage** - Controls whether cache devices are used for the share. ■ **Synchronous write bias** - Controls the behavior when servicing synchronous writes. By default, the system optimizes synchronous writes for latency, which leverages the log devices to provide fast response times. ■ **Database record size** - Controls the block size used by the filesystem. This property is only valid for filesystems. By default, filesystems will use a block size just large enough to hold the file, or 128K for large files. This means that any file over 128K in size will be using 128K blocks. If an application then writes to the file in small chunks, it will necessitate reading and writing out an entire 128K block, even if the amount of data being written is comparatively small. The property can be set to any power of 2 from 512 to 128K. ■ **Additional replication** - Controls number of copies stored of each block, above and beyond any redundancy of the storage pool. ■ **Virus scan** - Controls whether this filesystem is scanned for viruses. This property is only valid for filesystems. ■ **Prevent destruction** - When set, the share or project cannot be destroyed. This includes destroying a share through dependent clones, destroying a share within a project, or destroying a replication package.

Table 3–6 (Cont.) Project Settings

Section and Setting	Description
Default Settings	Custom settings for file systems, to be used as default, include the following: ■ **User** - User that is the current owner of the directory. ■ **Group** - Group that is the current owner of the directory. ■ **Permissions** - Permissions include Read (R), Write (W), or Execute (X). Ensure that you set the right permissions for users in a department such as `Dept_1`, that need write access to the shared file systems. Custom settings for LUNs, to be used as default, include the following: ■ **Volume Size** - Controls the size of the LUN. By default, LUNs reserve enough space to completely fill the volume ■ **Thin provisioned** - Controls whether space is reserved for the volume. This property is only valid for LUNs. By default, a LUN reserves exactly enough space to completely fill the volume. This ensures that clients will not get out-of-space errors at inopportune times. This property allows the volume size to exceed the amount of available space. When set, the LUN will consume only the space that has been written to the LUN. While this allows for thin provisioning of LUNs, most filesystems do not expect to get "out of space" from underlying devices, and if the share runs out of space, it may cause instability or ata corruption on clients, or both. ■ **Volume block size** - The native block size for LUNs. This can be any power of 2 from 512 bytes to 128K, and the default is 8K. For more information, see the Sun ZFS Storage documentation.

4. After entering your choices, click **Apply**.
5. After creating the `Dept_1` project, navigate to the `Dept_1` project page by clicking **PROJECTS**. Click **Protocols**. The project page is displayed, as in Figure 3–6.

Figure 3–6 Dept1_Project

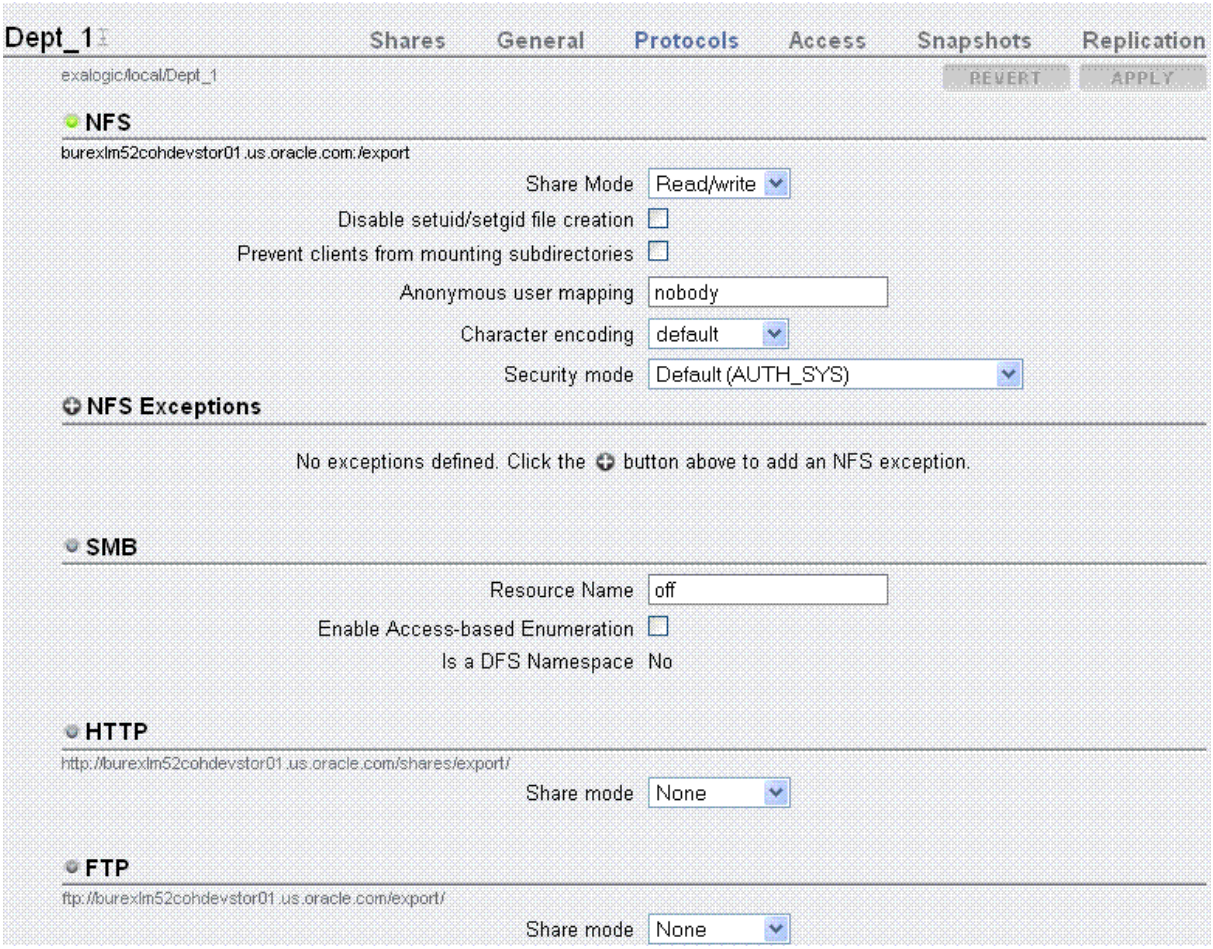

6. Click the **+** icon next to **NFS Exceptions** to add an NFS exception as follows:
   - Select **Network** as the **TYPE**.
   - In the **ENTITY** field, enter the IP address of the private subnet that you created for `Dept_1` spanning `ComputeNode1` and `ComputeNode2`.
   - Select the **ROOT ACCESS** option.
   - Click **APPLY**.

   Note: Using NFS Exceptions allows the `root` user of a compute node (operating system) to access Sun ZFS Storage 7320 appliance, which is defined for a particular user.

7. Repeat these steps to create a similar project named `FMW_Product1`.

3.4.2.2 Creating Shares for Dept_1

In the `Dept_1` project, create a separate share for Domain Home (`domains`) and a separate share for JMS persistence logs JTA transaction logs (`jmsjta`).

For example, to create the `domains` share for the `Dept_1` project, do the following:

1. In the Browser User Interface (BUI), access the Projects user interface by clicking **Configuration** > **STORAGE** > **Shares** > **Projects**. The Project Panel is displayed.
2. On the Project Panel, click `Dept_1`. The following page is displayed.

Figure 3–7 Dept_1 Project Page

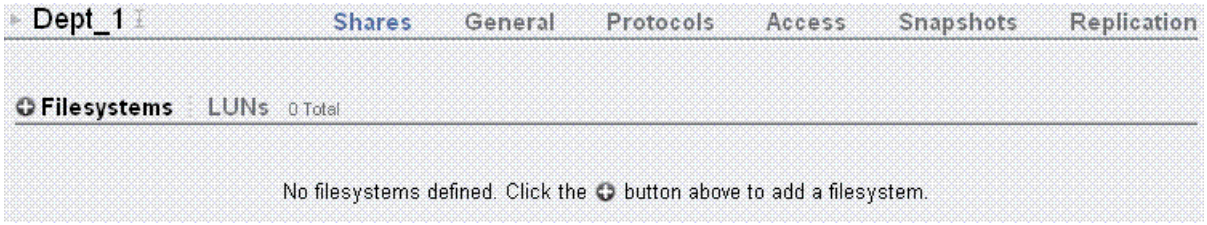

3. Click the **+** button next to **Filesystems** to add a file system. The Create Filesystem screen is displayed.

Figure 3–8 Create Filesystem

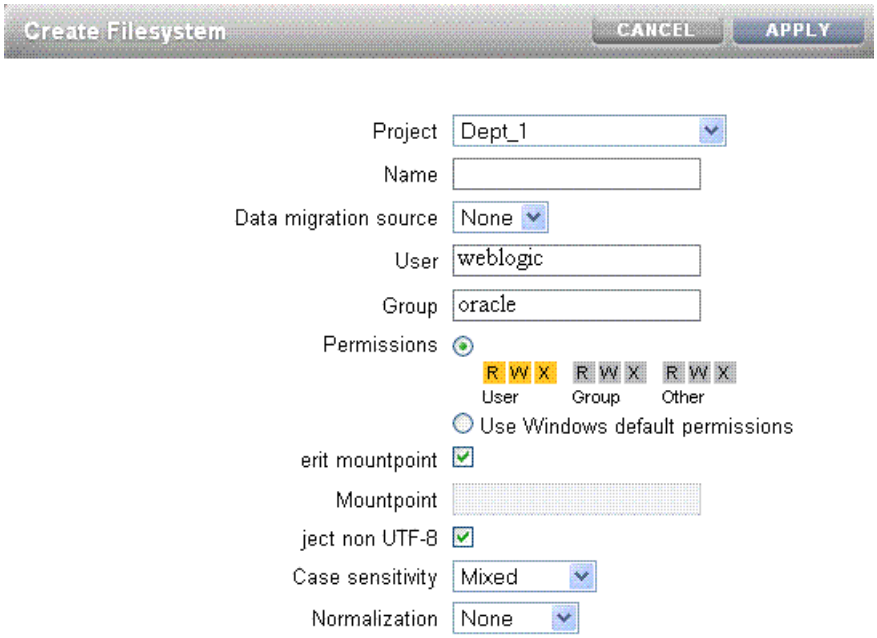

4. In the Create Filesystems screen, choose the target project from the **Project** pull-down menu. For example, choose `Dept_1`.
5. In the **Name** field, enter a name for the share. For example, enter `domains`.
6. From the **Data migration source** pull-down menu, choose **None**.
7. Select the **Permissions** option. Table 3–7 lists the access types and permissions.

Table 3–7 File System Access Types and Permissions

Access Type	Description	Permissions to Grant
User	User that is the current owner of the directory.	The following permissions can be granted: ■ **R** - Read - Permission to list the contents of the directory. ■ **W** - Write - Permission to create files in the directory. ■ **X** - Execute - Permission to look up entries in the directory. If users have execute permissions but not read permissions, they can access files explicitly by name but not list the contents of the directory. Ensure that the user `weblogic` has read-write permissions for the appropriate shares. For more information, see Section 3.4.2.9, "Creating Groups and Users and Controlling Access to Mounted Shares".
Group	Group that is the current group of the directory.	
Other	All other accesses.	

You can use this feature to control access to the file system, based on the access types (users and groups) in `Dept_1`.

8. You can either inherit mountpoint by selecting the **Inherit mountpoint** option or set a mountpoint.

   Note: The mount point must be under `/export`. The mount point for one share cannot conflict with another share. In addition, it cannot conflict with another share on cluster peer to allow for proper failover.

   When inheriting the mountpoint property, the current dataset name is appended to the project's mountpoint setting, joined with a slash ('/'). For example, if the `domains` share has the mountpoint setting `/export/domains`, then `domains/config` inherits the mountpoint `/export/domains/config`.

9. To enforce UTF-8 encoding for all files and directories in the file system, select the **Reject non UTF-8** option. When set, any attempts to create a file or directory with an invalid UTF-8 encoding will fail.

   Note: This option is selected only when you are creating the file system.

10. From the **Case sensitivity** pull-down menu, select **Mixed**, **Insensitive**, or **Sensitive** to control whether directory lookups are case-sensitive or case-insensitive.

Table 3–8 Case Sensitivity Values

BUI Value	Description
Mixed	Case sensitivity depends on the protocol being used. For NFS, FTP, and HTTP, lookups are case-sensitive. This is default, and prioritizes conformance of the various protocols over cross-protocol consistency.

Table 3–8 (Cont.) Case Sensitivity Values

BUI Value	Description
Insensitive	All lookups are case-insensitive, even over protocols (such as NFS) that are traditionally case-sensitive. This setting should only be used where CIFS is the primary protocol and alternative protocols are considered second-class, where conformance to expected standards is not an issue.
Sensitive	All lookups are case-sensitive. In general, do not use this setting.

Note: This option is selected only when you are creating the file system.

11. From the **Normalization** pull-down menu, select **None**, **Form C**, **Form D**, **Form KC**, or **Form KD** to control what unicode normalization, if any, is performed on filesystems and directories. Unicode supports the ability to have the same logical name represented by different encoding. Without normalization, the on-disk name stored will be different, and lookups using one of the alternative forms will fail depending on how the file was created and how it is accessed. If this property is set to anything other than **None** (the default), the **Reject non UTF-8** property must also be selected.

Table 3–9 Normalization Settings

BUI Value	Description
None	No normalization is done.
Form C	Normalization Form Canonical Composition (NFC) - Characters are decomposed and then recomposed by canonical equivalence.
Form D	Normalization Form Canonical Decomposition (NFD) - Characters are decomposed by canonical equivalence.
Form KC	Normalization Form Compatibility Composition (NFKC) - Characters are decomposed by compatibility equivalence, then recomposed by canonical equivalence.
Form KD	Normalization Form Compatibility Decomposition (NFKD) - Characters are decomposed by compatibility equivalence.

Note: This option is selected only when you are creating the file system.

12. After entering the values, click **Apply**. A share named `domains` is created and listed on the `Dept_1` project page.
13. Repeat these steps to create a similar share for `jmsjta` under the `Dept_1` project.

3.4.2.3 Creating Shares for FMW_Product1

You must also create shares for the `FMW_Product1` project, which can be accessed by all compute nodes. This project contains the shares for Oracle WebLogic product binaries, such as `wlserver_10.3`.

To create the `wlserver_10.3` share under `FMW_Product1`, use the steps described in Section 3.4.2.2, "Creating Shares for Dept_1" as a reference to create shares that `FMW_Product1` requires.

Additionally, you can create shares for the Oracle Enterprise Management Agent (master) if you optionally wish to use Oracle Enterprise Manager Grid Control to monitor software in your enterprise deployment topology.

3.4.2.4 NFSv4 Configuration Requirements

For information about configuring NFSv4 on Exalogic, see the chapter "Configuring NFS Version 4 (NFSv4) on Exalogic" in *Oracle Exalogic Machine Owner's Guide*.

3.4.2.5 Creating Mount Points on ComputeNode1 and ComputeNode2

On the command line, run the following commands as a `root` user on `ComputeNode1` and `ComputeNode2` to create the necessary mount points:

On ComputeNode1:

```
# mkdir -p /u01/common/patches
# mkdir -p /u01/common/general
# mkdir -p /u01/FMW_Product1/wlserver_1034
# mkdir -p /u01/FMW_Product1/webtier_1115
# mkdir -p /u01/Dept_1/domains/el01cn01/
# mkdir -p /u01/Dept_1/admin/
# mkdir -p /u01/Dept_1/jmsjta/
# mkdir -p /u01/el01cn01/dumps
# mkdir -p /u01/el01cn01/general
```

After creating these mount points, run the following commands:

```
# mkdir -p /u01/Dept_1/admin/el01cn01/nodemanager
# mkdir -p /u01/Dept_1/jmsjta/base_domain/Dept1_Cluster1/jms
# mkdir -p /u01/Dept_1/jmsjta/base_domain/Dept1_Cluster1/tlogs
```

On ComputeNode2:

```
# mkdir -p /u01/common/patches
# mkdir -p /u01/common/general
# mkdir -p /u01/FMW_Product1/wlserver_1034
# mkdir -p /u01/FMW_Product1/webtier_1115
# mkdir -p /u01/Dept_1/domains/el01cn02
# mkdir -p /u01/Dept_1/admin/
# mkdir -p /u01/Dept_1/jmsjta/
# mkdir -p /u01/el01cn02/dumps
# mkdir -p /u01/el01cn02/general
```

After creating these mount points, run the following commands:

```
# mkdir -p /u01/Dept_1/admin/el01cn02/nodemanager
# mkdir -p /u01/Dept_1/jmsjta/base_domain/Dept1_Cluster1/jms
# mkdir -p /u01/Dept_1/jmsjta/base_domain/Dept1_Cluster1/tlogs
```

Note: In the above examples, el01cn01 represents the host name assigned to ComputeNode1 and el01cn02 represents the host name assigned to ComputeNode2.

3.4.2.6 Editing the /etc/fstab (Linux) or /etc/vfstab (Solaris) File

After creating the mount points, you must add entries for the mount points to the /etc/fstab (Linux) or /etc/vfstab (Solaris) file on ComputeNode1 and ComputeNode2.

Log in as a root user, and complete the following steps:

1. On ComputeNode1, add the following entries to the /etc/fstab (Linux) or /etc/vfstab (Solaris) file in a text editor, such as vi:

   Oracle Linux

   - el01sn01-priv:/export/common/general /u01/common/general nfs4 rw,bg,hard,nointr,rsize=131072,wsize=131072
   - el01sn01-priv:/export/common/patches /u01/common/patches nfs4 rw,bg,hard,nointr,rsize=131072,wsize=131072
   - el01sn01-priv:/export/el01cn01/dumps /u01/el01cn01/dumps nfs4 rw,bg,hard,nointr,rsize=131072,wsize=131072
   - el01sn01-priv:/export/el01cn01/general /u01/el01cn01/general nfs4 rw,bg,hard,nointr,rsize=131072,wsize=131072
   - el01sn01-priv:/export/Dept_1/domains /u01/Dept_1/domains nfs4 rw,bg,hard,nointr,rsize=131072,wsize=131072
   - el01sn01-priv:/export/Dept_1/jmsjta /u01/Dept_1/jmsjta nfs4 rw,bg,hard,nointr,rsize=135268,wsize=135168
   - el01sn01-priv:/export/Dept_1/admin /u01/Dept_1/admin nfs4 rw,bg,hard,nointr,rsize=131072,wsize=131072
   - el01sn01-priv:/export/FMW_Product1/wlserver_1034 /u01/FMW_Product1/wlserver_1034 nfs4 rw,bg,hard,nointr,rsize=131072,wsize=131072
   - el01sn01-priv:/export/FMW_Product1/webtier_1115 /u01/FMW_Product1/webtier_1115 nfs4 rw,bg,hard,nointr,rsize=131072,wsize=131072

   Oracle Solaris

   - el01sn01-priv:/export/common/general - /u01/common/general nfs - yes rw,bg,hard,nointr,rsize=131072,wsize=131072,vers=4
   - el01sn01-priv:/export/common/patches - /u01/common/patches nfs - yes rw,bg,hard,nointr,rsize=131072,wsize=131072,vers=4
   - el01sn01-priv:/export/el01cn01/dumps - /u01/el01cn01/dumps nfs - yes rw,bg,hard,nointr,rsize=131072,wsize=131072,vers=4

- el01sn01-priv:/export/el01cn01/general - /u01/el01cn01/general nfs - yes rw,bg,hard,nointr,rsize=131072,wsize=131072,vers=4
- el01sn01-priv:/export/Dept_1/domains - /u01/Dept_1/domains nfs - yes rw,bg,hard,nointr,rsize=131072,wsize=131072,vers=4
- el01sn01-priv:/export/Dept_1/jmsjta - /u01/Dept_1/jmsjta nfs - yes rw,bg,hard,nointr,rsize=135268,wsize=135168,vers=4
- el01sn01-priv:/export/Dept_1/admin - /u01/Dept_1/admin nfs - yes rw,bg,hard,nointr,rsize=131072,wsize=131072,vers=4
- el01sn01-priv:/export/FMW_Product1/wlserver_1034 - /u01/FMW_Product1/wlserver_1034 nfs - yes rw,bg,hard,nointr,rsize=131072,wsize=131072,vers=4
- el01sn01-priv:/export/FMW_Product1/webtier_1115 - /u01/FMW_Product1/webtier_1115 nfs - yes rw,bg,hard,nointr,rsize=131072,wsize=131072,vers=4

Note: In the above entries, el01sn01-priv is used as the example host name of the Sun ZFS Storage 7320 appliance. You can also use the IPoIB IP address assigned to the storage appliance.

2. On ComputeNode2, add the following entries to the /etc/fstab (Linux) or /etc/vfstab (Solaris) file in a text editor, such as vi:

   Oracle Linux

   - el01sn01-priv:/export/common/general /u01/common/general nfs4 rw,bg,hard,nointr,rsize=131072,wsize=131072
   - el01sn01-priv:/export/common/patches /u01/common/patches nfs4 rw,bg,hard,nointr,rsize=131072,wsize=131072
   - el01sn01-priv:/export/el01cn02/dumps /u01/el01cn02/dumps nfs4 rw,bg,hard,nointr,rsize=131072,wsize=131072
   - el01sn01-priv:/export/el01cn02/general /u01/el01cn02/general nfs4 rw,bg,hard,nointr,rsize=131072,wsize=131072
   - el01sn01-priv:/export/FMW_Product1/wlserver_1034 /u01/FMW_Product1/wlserver_1034 nfs4 rw,bg,hard,nointr,rsize=131072,wsize=131072
   - el01sn01-priv:/export/FMW_Product1/webtier_1115 /u01/FMW_Product1/webtier_1115 nfs4 rw,bg,hard,nointr,rsize=131072,wsize=131072
   - el01sn01-priv:/export/Dept_1/jmsjta /u01/Dept_1/jmsjta nfs4 rw,bg,hard,nointr,rsize=135268,wsize=135168
   - el01sn01-priv:/export/Dept_1/admin /u01/Dept_1/admin nfs4 rw,bg,hard,nointr,rsize=131072,wsize=131072
   - el01sn01-priv:/export/Dept_1/domains /u01/Dept_1/domains nfs4 rw,bg,hard,nointr,rsize=131072,wsize=131072

Oracle Solaris

- el01sn01-priv:/export/common/general - /u01/common/general nfs - yes rw,bg,hard,nointr,rsize=131072,wsize=131072,proto=tcp,vers=4
- el01sn01-priv:/export/common/patches - /u01/common/patches nfs - yes rw,bg,hard,nointr,rsize=131072,wsize=131072,proto=tcp,vers=4
- el01sn01-priv:/export/el01cn02/dumps - /u01/el01cn02/dumps nfs - yes rw,bg,hard,nointr,rsize=131072,wsize=131072,proto=tcp,vers=4
- el01sn01-priv:/export/el01cn02/general - /u01/el01cn02/general nfs - yes rw,bg,hard,nointr,rsize=131072,wsize=131072,proto=tcp,vers=4
- el01sn01-priv:/export/FMW_Product1/wlserver_1034 - /u01/FMW_Product1/wlserver_1034 nfs - yes rw,bg,hard,nointr,rsize=131072,wsize=131072,proto=tcp,vers=4
- el01sn01-priv:/export/FMW_Product1/webtier_1115 - /u01/FMW_Product1/webtier_1115 nfs - yes rw,bg,hard,nointr,rsize=131072,wsize=131072,proto=tcp,vers=4
- el01sn01-priv:/export/Dept_1/jmsjta - /u01/Dept_1/jmsjta nfs - yes rw,bg,hard,nointr,rsize=135268,wsize=135168,proto=tcp,vers=4
- el01sn01-priv:/export/Dept_1/admin - /u01/Dept_1/admin nfs - yes rw,bg,hard,nointr,rsize=131072,wsize=131072,proto=tcp,vers=4
- el01sn01-priv:/export/Dept_1/domains - /u01/Dept_1/domains nfs - yes rw,bg,hard,nointr,rsize=131072,wsize=131072,proto=tcp,vers=4

Note: In the above entries, `el01sn01-priv` is used as the example host name of the Sun ZFS Storage 7320 appliance. You can also use the IPoIB IP address assigned to the storage appliance.

3. Save the file and exit.

3.4.2.7 Mounting the Volumes

To mount the volumes, complete the following steps:

1. On `ComputeNode1` and `ComputeNode2`, ensure that the mount entries are added to the `/etc/fstab` (Linux) or `/etc/vfstab` (Solaris) file correctly.

2. Run the `mount -a` command on both `ComputeNode1` and `ComputeNode2` to mount the volumes.

3.4.2.8 Creating Directories

After creating the mount points, run the following commands to create directories:

On ComputeNode1:

```
# mkdir -p /u01/Dept_1/admin/el01cn01/nodemanager
# mkdir -p /u01/Dept_1/jmsjta/base_domain/Dept1_Cluster1/jms
# mkdir -p /u01/Dept_1/jmsjta/base_domain/Dept1_Cluster1/tlogs
```

On ComputeNode2:

```
# mkdir -p /u01/Dept_1/admin/el01cn02/nodemanager
# mkdir -p /u01/Dept_1/jmsjta/base_domain/Dept1_Cluster1/jms
# mkdir -p /u01/Dept_1/jmsjta/base_domain/Dept1_Cluster1/tlogs
```

3.4.2.9 Creating Groups and Users and Controlling Access to Mounted Shares

After mounting the exported shares on the compute nodes, you can create groups and users on the operating system. From the Sun ZFS Storage 7320 appliance, you can set permissions for a particular share (exported file system) while creating a share.

For example, to create a primary group named `oinstall`, a secondary group named `oracle`, and a user name `weblogic` on `ComputeNode1`, run the following command on `ComputeNode1` as a `root` user:

```
# /usr/sbin/groupadd oinstall

# /usr/sbin/groupadd oracle

# /usr/sbin/useradd -g oinstall -G oracle weblogic

# su weblogic
```

You must change the ownership of the mount points to the `weblogic` user as follows:

```
# chown -R weblogic:oracle <MountPoint>
```

You must run the above command for each mountpoint. For more information about mount points, see Section 3.4.2.5, "Creating Mount Points on ComputeNode1 and ComputeNode2".

You will log in to the system as user `weblogic`.

Note: The user `weblogic` on the operating system should have the same user ID as the user that you created on the storage appliance while creating a new file system, as shown in Figure 3–8. You can use this user account `weblogic` to perform WebLogic Server product installation and configuration tasks. Do not perform these tasks as a `root` user. When setting up file systems in the BUI of the Sun ZFS Storage 7320 appliance, set the right permissions. The user `weblogic` should have read-write permissions to the relevant shares. For more information, see Section 3.4.2.2, "Creating Shares for Dept_1".

3.4.3 Recommendation About Storage Location for Syslogs and Operating System Patches

Syslogs are stored on the local Flash drives of Exalogic compute nodes. This option results in fast boot-up. However, you can configure log rotation to move syslogs to the Sun ZFS Storage 7320 appliance. For example, you can move syslogs of `ComputeNode1` to the `/u01/el01cn01/general` share on the Sun ZFS Storage 7320 appliance. `<node1>` represents the host name assigned to `ComputeNode1`.

Patches or updates required by the base operating system installed on the compute nodes may be stored on a share on the Sun ZFS Storage 7320 appliance.

3.5 Database

The Oracle Exalogic enterprise deployment reference topology uses the following database connectivity options:

- Oracle Exalogic machine connected to Oracle database or RAC over 10 Gb Ethernet

  You can connect an Oracle Real Application Clusters (RAC) database to an Oracle Exalogic machine on a 10 GB Ethernet link. This setup requires additional network configuration, such as creating a VNIC, as described in Section 3.5.2, "Connecting to Oracle Database Over Ethernet".

- Oracle Exalogic machine connected to Oracle Exadata Database Machine on the same InfiniBand fabric.

  For information about connecting Exalogic Machine to Oracle Exadata Database Machine, see *Oracle Fusion Middleware Exalogic Owner's Guide*.

The database preconfiguration, either using Oracle Exadata Database Machine or using a RAC database, is a prerequisite for configuring Oracle Fusion Middleware products and applications in the Oracle Exalogic environment.

This section covers the following topics:

- Section 3.5.1, "Prerequisite"
- Section 3.5.2, "Connecting to Oracle Database Over Ethernet"

3.5.1 Prerequisite

It is assumed that you have set up an Oracle RAC 11*g* Release 1 or 11*g* Release 2 database. This guide does not discuss how to create or set up a database.

3.5.2 Connecting to Oracle Database Over Ethernet

This section describes how to establish connectivity between an Oracle Exalogic machine and Oracle database (or RAC) over 10 Gb Ethernet. This scenario applies to Exalogic enterprise deployment scenarios where Oracle Exadata Database Machine is not present in the data tier.

This section contains the following sections:

- Section 3.5.2.1, "Overview"
- Section 3.5.2.2, "Prerequisites"
- Section 3.5.2.3, "Setting Up VNICs"

3.5.2.1 Overview

You can connect to Oracle Database via a dedicated physical LAN or a VLAN. However, both connectivity options require separate vNIC.

Use this section only if you want to connect your Oracle Exalogic machine to an Oracle database or RAC over a 10 Gb Ethernet link.

3.5.2.2 Prerequisites

Ensure that you have completed the following prerequisites before setting up network connectivity between the Oracle Exalogic machine and Oracle database over Ethernet:

- Ensure that IP addresses of all of the Sun Network QDR InfiniBand Gateway Switches (NM2-GW) in your Exalogic Machine are set up. Subnet Manager runs on these switches as a master or standby. To set up the IP address, run the following command:

  ```
  # echo -e
  \"aaa.aaa.aaa.aaa\\nbbb.bbb.bbb.bbb\\nccc.ccc.ccc.ccc\" >>
  /conf/smnodes
  ```

  Where `aaa.aaa.aaa.aaa`, `bbb.bbb.bbb.bbb`, and so on are the IP addresses of the Sun Network QDR InfiniBand Gateway Switches (NM2-GW) that run Subnet Manager either as master or as standby. The last IP address in the above command should not be followed by `\n`.

- Run Subnet Manager (Master) on one of the gateway switches. For more information, see the "Subnet Manager Requirements for Connecting Exalogic Machine to Oracle Exadata Database Machine" section in the *Oracle Fusion Middleware Exalogic Machine Owner's Guide*.

- Verify that the default InfiniBand partition at the Exalogic Machine level is configured. Verify the partition key. You can verify the default partition and the partition key by running the `smpartition list` command on the CLI of one of the gateway switches.

  At the command prompt on the CLI, enter the following command:

  ```
  # smpartition list active
  ```

  The command displays the active configuration of the InfiniBand partition, as in the following example:

  ```
  # smpartition list active
  # Sun DCS IB partition config file
  # This file is generated, do not edit
  #! version_number : 12
  Default=0x7fff, ipoib : ALL_CAS=full, ALL_SWITCHES=full, SELF=full;
  SUN_DCS=0x0001, ipoib : ALL_SWITCHES=full;
  #
  ```

  Note: For more information about these topics, see the *Oracle Fusion Middleware Exalogic Machine Owner's Guide*.

3.5.2.3 Setting Up VNICs

For information about creating a network interface on your Exalogic compute nodes for 10 GbE connectivity, see the "Configuring Ethernet Over InfiniBand" topic in the *Oracle Exalogic Machine Owner's Guide*.

4

Installing Oracle Software

This chapter describes the software installation required for the enterprise deployment reference topology.

This chapter contains the following sections:

- Section 4.1, "Installing Oracle WebLogic Software and Creating the Middleware Home"
- Section 4.2, "Setting Up Oracle Enterprise Manager Grid Control"
- Section 4.3, "Installing Oracle HTTP Server"

4.1 Installing Oracle WebLogic Software and Creating the Middleware Home

You must complete the following steps:

- Section 4.1.1, "Downloading the Oracle WebLogic Software Installer"
- Section 4.1.2, "Installing JDKs on Oracle Solaris"
- Section 4.1.3, "Installing Oracle WebLogic Server and Creating Middleware Home on Sun ZFS Storage 7320 appliance"
- Section 4.1.4, "Backing Up Installation"

4.1.1 Downloading the Oracle WebLogic Software Installer

The Oracle WebLogic Server 11*g* Release 1 (10.3.4) Installer includes the following components:

- Oracle WebLogic Server 11*g* Release 1 (10.3.4)
- Oracle Coherence 3.6.0.4
- Oracle JRockit JVM 28.1.0 (Oracle Linux)
- Sun JDK 1.6.0_23 (Oracle Solaris)

You must download the **Oracle WebLogic Linux (64-bit JVM)** or **Oracle WebLogic Solaris (64-bit JVM)** installer as follows:

1. Download the ZIP file from the Oracle E-Delivery Web site at:

   `http://edelivery.oracle.com`

2. Copy the `wls1034_linux64` or `wls1034_solaris64` to a local directory on `/u01/common/general`.

3. Extract the contents of the zip file to a local directory on `/u01/common/general`.

> **Note:** You can download the Oracle WebLogic 10.3.4 software from `http://edelivery.oracle.com`. Select **Oracle Fusion Middleware** as the Product Pack, **Linux x86-64** or **Oracle Solaris on x86-64 (64-bit)** as the Platform, and **Oracle Fusion Middleware 11g Media Pack for Exalogic** as the Media Pack.

4.1.2 Installing JDKs on Oracle Solaris

You must install the JDKs on Oracle Solaris as follows:

1. Download Oracle JDK 6 Update 25 for Oracle Solaris on x86 (32-bit) and Oracle JDK 6 Update 25 for Oracle Solaris on x86-64 (64-bit) included in the Oracle Exalogic Elastic Cloud Software 11*g* Media Pack available at:

   `http://edelivery.oracle.com`

2. Install Oracle JDK 6 Update 25 for Oracle Solaris on x86 (32-bit), as described in `http://www.oracle.com/technetwork/java/javase/install-solaris-139361.html#install`.

3. Install Oracle JDK 6 Update 25 for Oracle Solaris on x86-64 (64-bit), as described in `http://www.oracle.com/technetwork/java/javase/install-solaris-64-138849.html`.

4.1.3 Installing Oracle WebLogic Server and Creating Middleware Home on Sun ZFS Storage 7320 appliance

As described in Section 3.4, "Shared Storage and Recommended Project and Share Structure," you install the Oracle WebLogic product binaries on one of the shares in the Sun ZFS Storage 7320 appliance. Note that the share, which is a shared file system, must be accessible by all compute nodes.

You must run the Oracle WebLogic installer on `ComputeNode1` as follows:

1. Open a Linux or Solaris terminal window.

> **Note:** Ensure that you do not run the Oracle WebLogic installer as the `root` user for Linux. You can log in as user `weblogic` with write privileges to the appropriate shares mounted on the Sun ZFS Storage 7320 appliance. For more information, see Section 3.4.2.9, "Creating Groups and Users and Controlling Access to Mounted Shares".

2. Use the `cd` command to move from your present working directory to the directory where you extracted the contents of the zip file to.

3. At the command prompt, run the following command to start the installer.

   Oracle Linux:

   ```
   ./wls1034_linux64.bin
   ```

   Oracle Solaris:

   ```
   $JAVA_HOME/jdk1.6.0_25/bin/java -d64 -jar wls1034_generic.jar
   ```

   The **Welcome** screen is displayed.

4. In the **Welcome** screen, click **Next**.

   The **Choose Middleware Home Directory** screen is displayed.

5. In the Choose Middleware Home Directory screen, do the following:

   - Select **Create a New Middleware Home**.
   - Click **Browse**, and specify the shared file system on the Sun ZFS Storage 7320 appliance. For example, `/u01/app/FMW_Product1/Oracle/Middleware`.

   Note: This shared file system must be accessible by all compute nodes in the Oracle Exalogic machine.

   Click **Next**.

   The **Register for Security Updates** screen is displayed.

6. In the **Register for Security Updates** screen, enter your contact information, so that you can be notified of security updates, and click **Next**.

   The **Choose Install Type** screen is displayed.

7. In the **Choose Install Type** screen, select **Custom**, and click **Next**.

   The **Choose Products and Components** screen is displayed.

8. In the **Choose Products and Components** screen, leave the default values, but ensure that the following components are not selected:

   - **Server Examples** under **WebLogic Server**
   - **Coherence Examples** under **Oracle Coherance**

   Click **Next**.

   The **JDK Selection** screen is displayed.

9. In the **JDK Selection** screen, select **Oracle JRockit 1.6.0_20 SDK** or **Sun JDK 1.6.0_23** (Oracle Solaris), and click **Next**.

   The **Choose Product Installation Directories** screen is displayed.

10. In the **Choose Product Installation Directories** screen, ensure that the path to the shared file system where you want to copy the installation files is specified, and click **Next**.

    The **Installation Summary** screen is displayed.

11. In the **Installation Summary** screen, click **Next**.

    The **Installation Complete** screen is displayed.

12. In the Installation Complete screen, de-select **Run QuickStart**, and click **Done** to exit the installer.

For more information about Oracle WebLogic Installation types, log files, silent installation, and troubleshooting, see the *Oracle Fusion Middleware Installation Guide for Oracle WebLogic Server*.

4.1.4 Backing Up Installation

You must back up the Oracle Fusion Middleware Home created on the shared file system on the Sun ZFS Storage 7320 appliance:

```
ComputeNode1> tar -cvzf /u01/app/Dept_1/backup/fmwhomeback.tgz /u01/app/FMW_
```

```
Product1/Oracle/Middleware
```

This command creates a backup of the installation files for Oracle WebLogic.

4.2 Setting Up Oracle Enterprise Manager Grid Control

This step is optional, but recommended.

Oracle recommends that you set up and use Oracle Enterprise Manager to monitor software components, including Oracle WebLogic domains, application deployments, Coherence Clusters, and hosts.

Oracle Enterprise Manager Grid Control requires an Oracle Enterprise Management Agent and the Oracle Management Service (OMS). Oracle Management Service is outside of Oracle Exalogic environment. A single Oracle Enterprise Management Agent, required as Master Agent, is installed on the Sun ZFS Storage 7320 appliance. This master agent, which is on the shared storage is used by all compute nodes in the Oracle Exalogic machine.

You must install Oracle Management NFS agents on each of the local storage of the compute nodes.

Setting up Oracle Enterprise Manager in the Oracle Exalogic environment involves the following steps:

1. Patching the Oracle Management Service 11*g* repository, installed outside of the Exalogic environment. This patch includes Exalogic-specific features and enhancements. For information about patch requirements, see the Knowledge Base article (Title: *Oracle Exalogic Elastic Cloud 11g R1 - Known Issues* Doc ID: **1268557.1**) located at *My Oracle Support*.
2. Installing Enterprise Manager System on a shared file system on the Sun ZFS Storage 7320 appliance. For more information, see "*Oracle Enterprise Manager Grid Control Advanced Installation and Configuration Guide*".
3. Installing Oracle Management Agent on a shared file system on the Sun ZFS Storage 7320 appliance. This single agent will be used as the Master Agent by all Exalogic compute nodes.

   For information, see "Installing Oracle Management Agent Using Shared Oracle Home" in the *Oracle Enterprise Manager Grid Control Advanced Installation and Configuration Guide*.
4. Patching the Oracle Enterprise Management Agent (Master Agent) installed on the shared file system on the Sun ZFS Storage 7320 appliance. This patch includes Exalogic-specific features and enhancements. For information about patch requirements, see the Knowledge Base article (Title: *Oracle Exalogic Elastic Cloud 11g R1 - Known Issues* Doc ID: **1268557.1**) located at *My Oracle Support*.
5. Installing Oracle Management NFS Agents on Exalogic compute nodes. For more information, see "Installing Oracle Management Agent Using Shared Oracle Home Using nfsagentinstall Script" in the *Oracle Enterprise Manager Grid Control Advanced Installation and Configuration Guide*.

4.3 Installing Oracle HTTP Server

In the enterprise deployment reference topology described in this guide, Oracle HTTP Server is used as follows:

- Mandatorily in the web tier in the DMZ public zone to route requests from Oracle HTTP Server to WebLogic Server in the application tier - this installation of Oracle HTTP Server is external to the Exalogic machine.

  Note: For information about installing Oracle HTTP Server 11*g* Release1 (11.1.1.5.0) outside of the Exalogic environment, see *Oracle Fusion Middleware Installation Guide for Oracle Web Tier*.

  You must configure Oracle HTTP Server, as described in <link>

  In the configuration example used in this guide, `WEBHOST1` and `WEBHOST2` are used as the Web server instances supporting the WebLogic servers in the Exalogic environment. These instances route requests from Oracle HTTP Server to Oracle WebLogic Server.

- Optionally inside the Exalogic box to load balance traffic on Exalogic's `BOND0/IPMP0` (IPoIB) network - this installation of Oracle HTTP Server resides on the shared storage appliance (Sun ZFS Storage 7320 appliance), and the instances run on Exalogic compute nodes.

  As described in Section 3.4, "Shared Storage and Recommended Project and Share Structure," you install the Oracle HTTP Server product binaries on one of the shares in the Sun ZFS Storage 7320 appliance. Note that the share, which is a shared file system, must be accessible by all compute nodes.

  Installing Oracle HTTP Server on Exalogic

  You must download the following Oracle HTTP Server 11*g* Release 1 installers to the `/u01/el01cn01/general` location:

  – Oracle HTTP Server 11*g* Release 1 (11.1.1.2.0) Installer

  – Oracle HTTP Server 11*g* Release 1 (11.1.1.5.0) Patch Set Installer

  See "Obtain the Oracle Fusion Middleware Software" in *Oracle Fusion Middleware Installation Planning Guide* for information on where to obtain the software.

  Select one of the download locations and download Oracle Web Tier 11.1.1.2.0 and 11.1.1.5.0 to the `/u01/el01cn01/general` location. These installers will be saved as .zip archive files.

  After you download the archive files, unpack the archive files into `/u01/el01cn01/general` on `ComputeNode1`.

  To install Oracle HTTP Server 11*g* Release 1 (11.1.1.5.0) on the Sun ZFS Storage 7320 appliance, complete the following steps:

  1. Start the installer, go to the directory where you unpacked the 11.1.1.2.0 archive file and switch to the `Disk1` directory.
  2. Run the following command to start the Oracle HTTP Server 11*g* Release 1 (11.1.1.2.0) installation program on `ComputeNode1`:

     ```
     ./runInstaller
     ```

     The Welcome screen of the Oracle Web Tier 11.1.1.2.0 installation wizard is displayed.
  3. On the Welcome screen, click **Next**. The Select Installation Type screen is displayed.

4. Select the **Install Software - Do Not Configure** option. Click **Next** to continue. The Prerequisite Checks screen is displayed.
5. When the prerequisite check is complete, click **Next** to continue. The Specify Installation screen is displayed.
6. Specify the Middleware Home (`/u01/FMW_Product1/Oracle/Middleware`) and Oracle Web Tier home (`/u01/FMW_Product1/Oracle/Middleware/Oracle_WT1`) directories. The Web Tier Oracle home and Oracle common home directories will be created inside the Middleware Home directory. Click **Next** to continue. The Specify Security Updates screen is displayed.
7. If you choose to register for security updates, provide your email address to be informed of the latest product issues. Click **Next** to continue. The Installation Summary screen is displayed.
8. Review the summary, and click **Install** to begin the installation. The Installation Progress screen is displayed. Click **Next** to continue. When the installation is complete, the Installation Complete screen is displayed.
9. Click **Finish** to dismiss the installer. The Web Tier 11.1.1.2.0 software, which includes Oracle HTTP Server 11.1.1.2.0, is installed.
10. Start the Oracle Web Tier 11.1.1.50 Patch Set installer, go to the directory where you unpacked the 11.1.1.5.0 archive file, and switch to the `Disk1` directory.
11. Run the following command to start the Oracle HTTP Server 11*g* Release 1 (11.1.1.5.0) installation program on `ComputeNode1`:

    ```
    ./runInstaller
    ```

    The Welcome screen of the Oracle Web Tier 11.1.1.5.0 patch set installation wizard is displayed.
12. Click **Next** to continue. The Specify Installation Location screen is displayed.
13. Specify your existing Middleware Home and Oracle Web Tier 11.1.1.2.0 home directories. Click **Next** to continue. The Specify Security Updates screen is displayed.
14. Enter your E-mail address if you want to receive the latest product information and security updates. If you have a My Oracle account and wish to receive updates via this mechanism, select **I wish to receive security updates via My Oracle Support**, then enter your account password. If you do not wish to register for security updates, leave all fields on this screen blank. Click **Next** to continue. The Installation Summary screen is displayed.
15. Review the installation summary, and click **Install** to continue. The Installation Progress screen is displayed. Click **Next** to continue. When the installation is complete, the Installation Complete screen is displayed.
16. Click **Finish** to dismiss the installer. Your Oracle Web Tier 11.1.1.2.0 installation is now updated to Oracle Web Tier 11.1.1.5.0.

Note: After installing Oracle HTTP Server 11.1.1.5.0, you must configure Oracle HTTP Server, as described in Section 6.4, "Optional: Configuring Oracle HTTP Server for Load Balancing on the Private InfiniBand Network".

Installing a Mandatory Patch for Oracle HTTP Server Installed on Exalogic

If you have installed Oracle HTTP Server to optionally load balance traffic on the private InfiniBand network in the Exalogic machine, you must install a mandatory patch before configuring Oracle HTTP Server instances to run on Exalogic compute nodes.

To install the patch, complete the following steps:

1. Go to **My Oracle Support** at `http://support.oracle.com`, click on the **Patches & Updates** tab, and search for the following patch number:

   For Oracle Linux and Oracle Solaris: **12530765**

2. Download the associated patch and install it by following the instructions in the README file included with the patch.

5

Configuring Oracle Fusion Middleware

After installing the Oracle WebLogic software, you must configure Oracle Fusion Middleware for the Oracle Exalogic enterprise deployment topology described in Chapter 2, "Reference Topology and Slicing Diagram". This configuration uses the following Oracle Fusion Middleware components:

- Oracle WebLogic Server 11*g* Release 1 (10.3.4)
- Oracle Coherence 3.6.0.4
- Oracle JRockit JVM 28.1.0 (For Oracle Linux)
- Oracle Sun JDK 1.6.0_23 (For Oracle Solaris)
- Oracle HTTP Server 11*g* Release 1 (11.1.1.4.0)

> **Note:** In this guide, Oracle HTTP Server is used as the web server in the Web tier of the enterprise deployment reference topology.

This chapter shows how to configure Oracle Fusion Middleware for two machines (referred to as Exalogic compute nodes) `ComputeNode1` and `ComputeNode2`, as illustrated in Figure 5–1.

You can follow this example configuration to create WebLogic domains with clusters of Managed Servers on the remaining compute nodes in the Oracle Exalogic machine, based on your application deployment and management requirements.

Figure 5–1 Example Configuration Scenario for Exalogic x86 Physical Machines

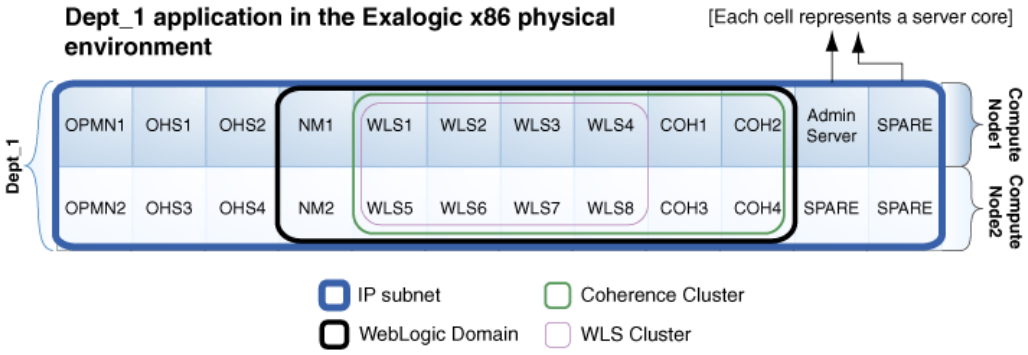

In this example configuration, you are creating a single Oracle WebLogic domain including the following:

- Four Managed Servers each on `ComputeNode1` and on `ComputeNode2` in a WebLogic cluster
- Node Manager on each compute node
- Two optional Oracle HTTP Server instances per compute node for load balancing traffic on the `BOND0` (IPoIB) network
- An OPMN process on each compute node to monitor the optional Oracle HTTP Server instances
- Oracle WebLogic Administration Server running on one of the compute nodes, such as `ComputeNode1`
- An Oracle Coherence cluster comprising four storage-enabled Coherence Servers spanning `ComputeNode1` and `ComputeNode2`

Note: In Figure 5–1, `Coh1` and `Coh2` represent Coherence nodes that are configured to run on `ComputeNode1`. `Coh3` and `Coh4` run on `ComputeNode2`. These are used as example names only. The nodes have their own end point (`BOND0` IP addresses of the machines as the host addresses). In the example scenario discussed in this guide, the `BOND0` IP of `ComputeNode1` is `192.168.10.1`. `ComputeNode2` uses `192.168.10.2`.

This chapter discusses the following topics:

5.1 Important Notes Before You Begin

Read the following notes before you start configuring Oracle Fusion Middleware components:

- If you are an Oracle Solaris user, read Section 3.1, "Important Notes for Oracle Solaris Users" before you configure Oracle Fusion Middleware.
- The configuration example used in this guide describes how to configure the environment for one department that uses two compute nodes (`Dept_1` using `ComputeNode1` and `ComputeNode2`). In this example, WebLogic Managed Servers run on `ComputeNode1` and `ComputeNode2`, the Administration Server runs on `ComputeNode1`, and separate Node Manager instances run on `ComputeNode1` and `ComputeNode2`.
- You can extrapolate the information included in these procedures to set up and configure the environment for your remaining departments, as necessary. This configuration depends on your specific application deployment and management requirements as well as the Oracle Exalogic machine configuration.
- Oracle Exalogic machine full rack includes 30 compute nodes, an Oracle Exalogic machine half rack includes 16 compute nodes, and Oracle Exalogic machine quarter rack includes 8 compute nodes. You should plan your application deployment and infrastructure accordingly.

5.2 Prerequisites

The following are the prerequisites for configuring Oracle Fusion Middleware 11*g* Release 1 (11.1.1) products for Oracle Exalogic:

- Preconfiguring the environment, including database, storage, and network, as described in Chapter 3, "Network, Storage, and Database Preconfiguration".
- Installing Oracle WebLogic Server 11*g* Release 1 (10.3.4) and creating a Middleware Home on a shared file system on Sun ZFS Storage 7320 appliance. This file system should be accessible by both `ComputeNode1` and `ComputeNode2`, as described in Section 4.1, "Installing Oracle WebLogic Software and Creating the Middleware Home".

5.3 Enabling Floating IP for Administration Server on ComputeNode1

The Administration Server must be configured to listen on a floating IP Address to enable it to seamlessly failover from one host to another. In case of a failure, the Administration Server, along with the virtual IP Address, can be migrated from one host to another.

You are associating the Administration Server with a virtual hostname (`ADMINVHN1`). This Virtual Host Name must be mapped to the appropriate floating IP (`10.0.0.17`) by a custom `/etc/hosts` entry. Check that the floating IP is available per your name

resolution system, (/etc/hosts), in the required nodes in your enterprise deployment reference topology. The floating IP (10.0.0.17) that is associated with this Virtual Host Name (ADMINVHN1) must be enabled on ComputeNode1.

To enable the floating IP on ComputeNode1, complete the following steps:

1. On ComputeNode1, run the ifconfig command as the root user to get the value of the netmask.

   For example:

   ```
   [root@ComputeNode1 ~] # /sbin/ifconfig
   bond0     Link encap:Ethernet  HWaddr 00:11:43:D7:5B:06
        inet addr:139.185.140.51  Bcast:139.185.140.255  Mask:255.255.255.224
        inet6 addr: fe80::211:43ff:fed7:5b06/64 Scope:Link
        UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
        RX packets:10626133 errors:0 dropped:0 overruns:0 frame:0
        TX packets:10951629 errors:0 dropped:0 overruns:0 carrier:0
        collisions:0 txqueuelen:1000
        RX bytes:4036851474 (3.7 GiB)  TX bytes:2770209798 (2.5 GiB)
        Base address:0xecc0 Memory:dfae0000-dfb00000
   ```

2. On ComputeNode1, bind the floating IP Address to the network interface card using ifconfig command as the root user. Use a netmask value that was obtained in Step 1.

   Note: For Oracle Solaris, you must plumb the interface: /sbin/ifconfig networkCardInterface plumb

   ```
   /sbin/ifconfig networkCardInterface Virtual_IP_Address netmask netMask
   ```

   For example, on Oracle Linux:

   ```
   /sbin/ifconfig bond0:1 10.0.0.17 netmask 255.255.255.224
   ```

   In this example, bond0:1 is the virtual network interface created for internal, fabric-level InfiniBand traffic.

   For example, on Oracle Solaris:

   ```
   /sbin/ifconfig ipmp0:1 plumb
   /sbin/ifconfig ipmp0:1 10.0.0.17 netmask 255.255.255.224 up
   ```

3. For Oracle Linux, run the arping command as the root user to update the routing table:

   ```
   /sbin/arping -q -U -c 3 -I networkCardInterface Floating_IP_Address
   ```

   For example:

   ```
   /sbin/arping -q -U -c 3 -I bond0 10.0.0.17
   ```

   Note: It is recommended that you run this command several times to update the routing table.

4. Verify the floating IP address of the Administration Server by running the netstat -nr command on ComputeNode1.

Note: In this enterprise deployment topology, example IP addresses are used. You must replace them with your own IP addresses that you reconfigured using `Exalogic Configuration Utility`. Even if the Administration Server for your WebLogic domain does not require a floating IP, it is recommended that you assign a floating IP address if you want to migrate the Administration Server manually from `ComputeNode1` to `ComputeNode2`.

5.4 Running Oracle Fusion Middleware Configuration Wizard on ComputeNode1 to Create an Oracle WebLogic Domain

Run the Oracle Fusion Middleware Configuration Wizard to create a Oracle WebLogic domain as follows:

Note: Ensure that you do not run the Oracle WebLogic installer as the `root` user. You can log in as user `weblogic` with write privileges to the appropriate shares mounted on the Sun ZFS Storage 7320 appliance. For more information, see Section 3.4.2.9, "Creating Groups and Users and Controlling Access to Mounted Shares".

1. On `ComputeNode1`, open a terminal window. At the command prompt, use the `cd` command to move from your present working directory:

   ```
   ComputeNode1> cd /u01/app/FMW_Product1/Oracle/Middleware/wlserver_
   10.3/common/bin
   ```

2. Start the Configuration Wizard:

   ```
   ComputeNode1> ./config.sh
   ```

   The **Welcome** screen is displayed.

3. Select **Create a new WebLogic Domain**, and click **Next**.

   The **Select Domain Source** screen is displayed.

4. Select the following components:

   - **Basic WebLogic Server Domain - 10.3.4.0 [wlserver_10.3]** (this should be selected automatically)
   - **Oracle Enterprise Manager - 11.1.1.0 [oracle_common]**
   - **Oracle JRF - 11.1.1.0 [oracle_common]**

   Click **Next**.

   The **Specify Domain Name and Location** screen is displayed.

5. Enter the following:

   - **Domain name**: Enter `base_domain` as the name for the Oracle WebLogic domain. This is the default name.
   - **Domain location**: Click **Browse** and specify the path `/u01/Dept_1/domains/el01cn01` in the shared file system on the Sun ZFS Storage 7320 appliance.

Ensure that the domain directory matches the directory and shared storage mount point recommended in Chapter 3, "Network, Storage, and Database Preconfiguration".

Click **Next**.

The **Configure Administrator User Name and Password** screen is displayed.

6. In this screen, enter the username and password to be used for the domain's administrator.

   Click **Next**.

   The **Configure Server Start Mode and JDK** screen is displayed.

7. Select the following:

   - From **WebLogic Domain Startup Mode**, select **Production Mode**.
   - From **Available JDKs**, select **JROCKIT SDK 1.6.0_20** (Oracle Linux) or **Sun JDK 1.6.0_23** (Oracle Solaris).

   Click **Next**.

   The **Select Optional Configuration** screen is displayed

8. Select the following:

   - **Administration Server**
   - **Managed Servers, Clusters and Machines**

   Click **Next**.

   The **Configure the Administration Server** screen is displayed.

9. Enter the following values:

   - Name: **AdminServer**
   - Listen address: **10.0.0.17**

     Note: This address, which is associated with the `BOND0` interface, is the example floating IP address assigned to the Administration Server

   - Listen port: **7001**
   - SSL listen port: **N/A**
   - SSL enabled: leave this check box unselected.

   Click **Next**.

   The **Configure Managed Servers** screen is displayed.

10. Click **Add**, and create eight Managed Servers each for `ComputeNode1` and `ComputeNode2` as shown in Table 5–2 and Table 5–2. In the **Listen address** enter the floating IP addresses of the Managed Servers (see Figure 5–2).

    Note: In this example, IP addresses are used as listen addresses. However, you can specify host names if they resolve to their corresponding floating IP addresses.

Table 5–1 Managed Servers on ComputeNode1

Name	Listen Address Using the BOND0 Interface via IPoIB
WLS1	10.0.0.1
WLS2	10.0.0.2
WLS3	10.0.0.3
WLS4	10.0.0.4

Table 5–2 Managed Servers on ComputeNode2

Name	Listen Address Using the BOND0 Interface via IPoIB
WLS5	10.0.0.5
WLS6	10.0.0.6
WLS7	10.0.0.7
WLS8	10.0.0.8

Tip: While creating Managed Servers, ensure that the names of the servers are unique across Oracle WebLogic domains in the Oracle Exalogic environment.

If you want to interoperate two or more domains via JMS or JTA, then you must set unique names for Oracle WebLogic domains, Oracle WebLogic Server, and JMS servers even if they are in different domains.

Figure 5–2 Configure Managed Servers

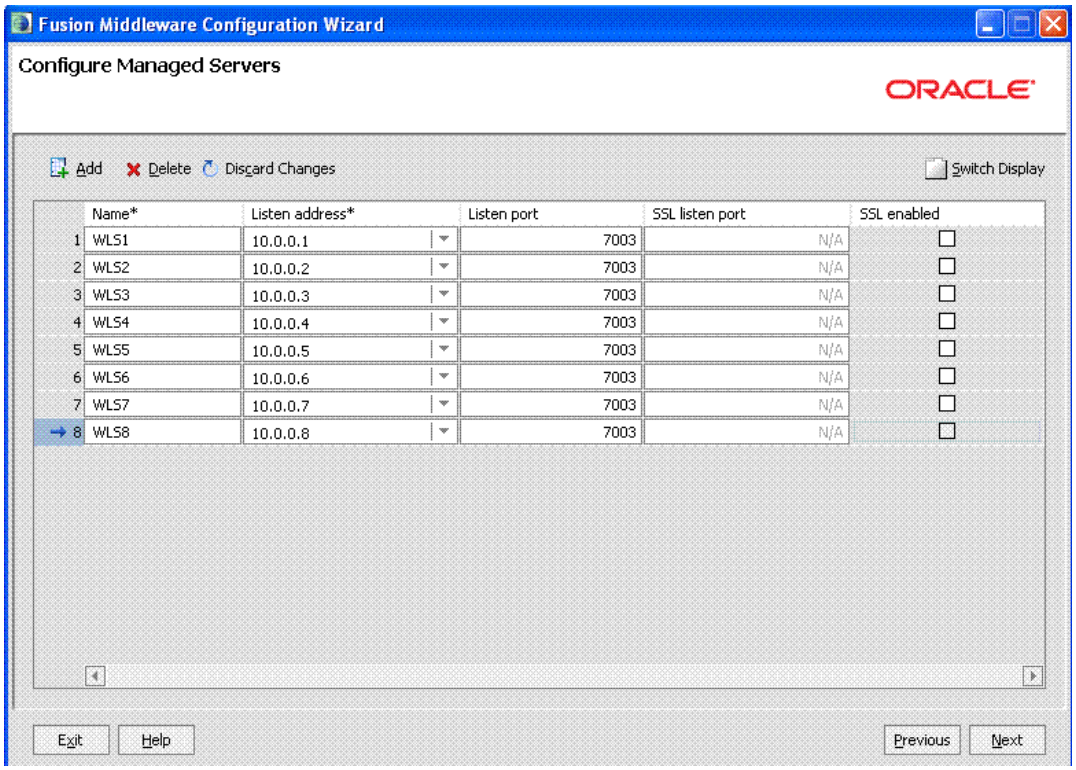

Note: In this enterprise deployment topology, floating IP addresses using the `BOND0` interface are assigned to Managed Servers and to the Administration Server. You must replace these IP addresses with your own IP addresses that you reconfigured using `Oracle Exalogic Configuration Utility`.

Click **Next**.

The **Configure Clusters** screen is displayed.

11. Click **Add** to configure clusters of Managed Servers on `ComputeNode1` and `ComputeNode2`.

    Enter a name for the new cluster, such as `Dept1_Cluster1` (see Figure 5–3).

Figure 5–3 Configure Clusters

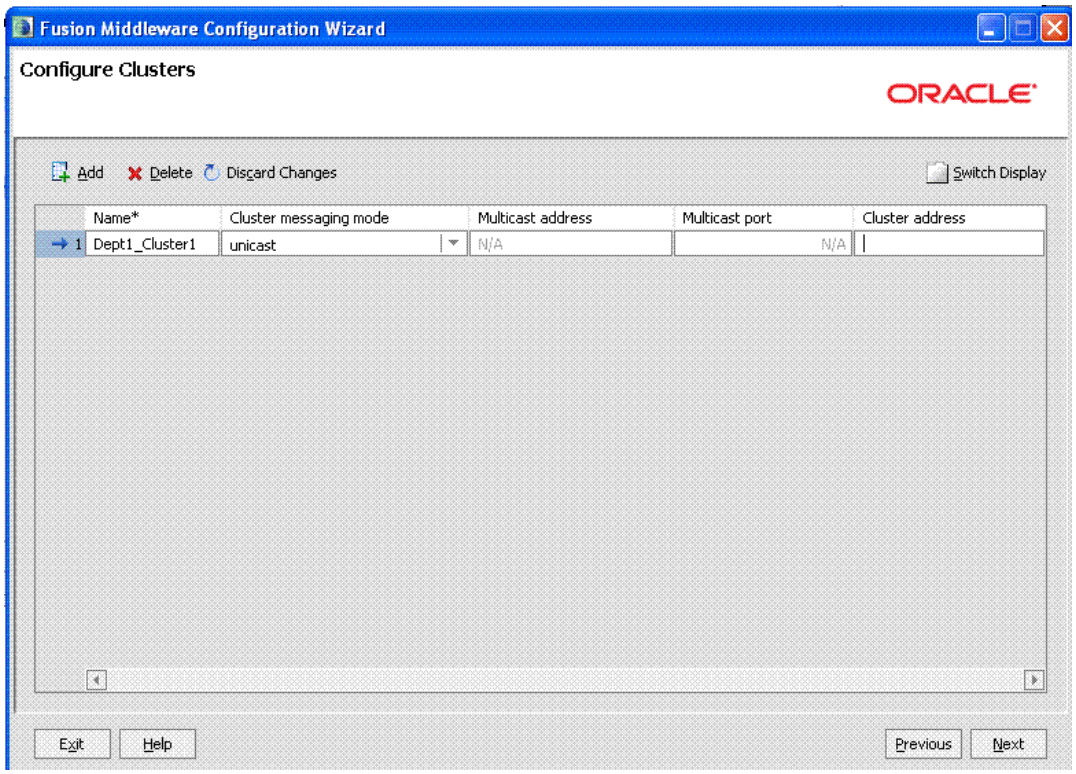

Click **Next**.

12. The **Assign Servers to Clusters** screen is displayed. In this screen, do the following:

    - Select **Dept1_Cluster1** in the **Cluster** pane, and double-click the following Managed Servers you created on `ComputeNode1` and `ComputeNode2`:
      - `WLS1`
      - `WLS2`
      - `WLS3`
      - `WLS4`

- WLS5
- WLS6
- WLS7
- WLS8

The names of the Managed Servers are removed from the **Server** pane and added below **Dept1_Cluster1** in the **Cluster** pane (see Figure 5–4).

Figure 5–4 Managed Servers Moved to Dept1_Cluster1

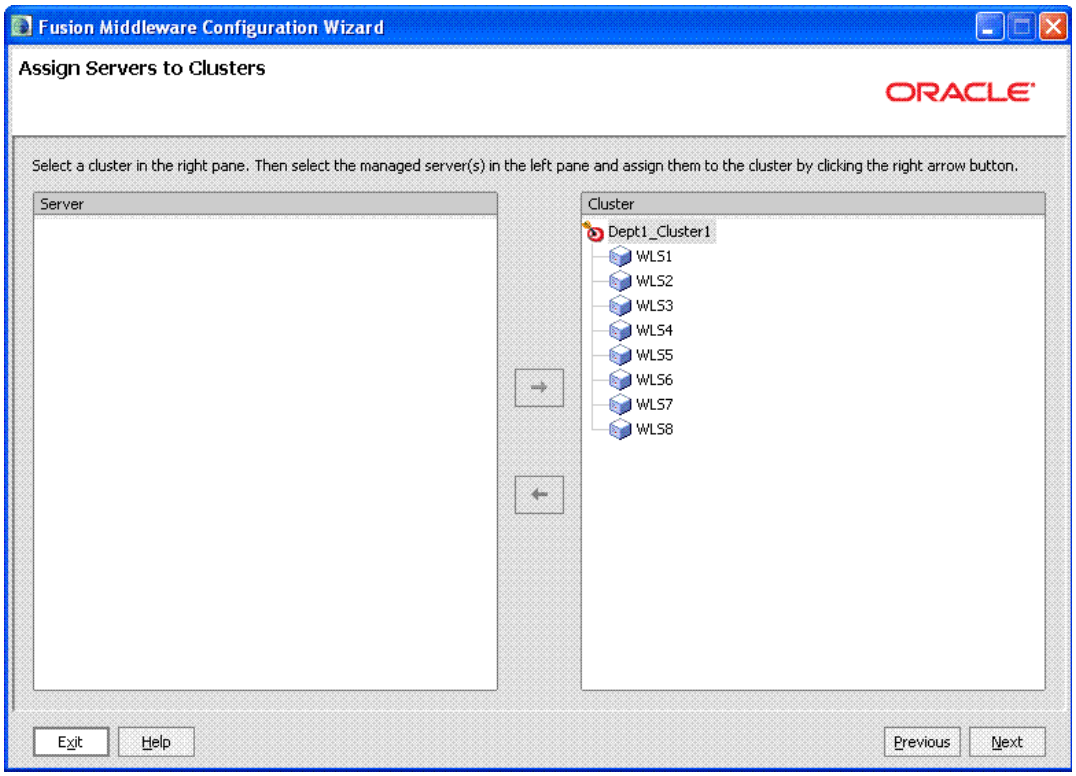

- Click **Next**.

The **Configure Machines** screen is displayed.

13. The **Configure Machines** screen is displayed. In this screen, click the **Unix Machine** tab and then click **Add** to add the following machines:

Table 5–3 Machines

Name	Node Manager Listen Address
ComputeNode1	el01cn01-priv The BOND0 IP address of ComputeNode1 is 192.168.10.1.
ComputeNode2	el01cn02-priv The BOND0 IP address of ComputeNode2 is 192.168.10.2.
ADMINHOST	ADMINVHN1 The floating IP address of ADMINHOST is 10.0.0.17.

Leave all other fields to their default values. Please note that the machine name does not need to be a valid host name or listen address; it is just a unique identifier of a Node Manager location (see Figure 5–5).

Figure 5–5 Configure Machines

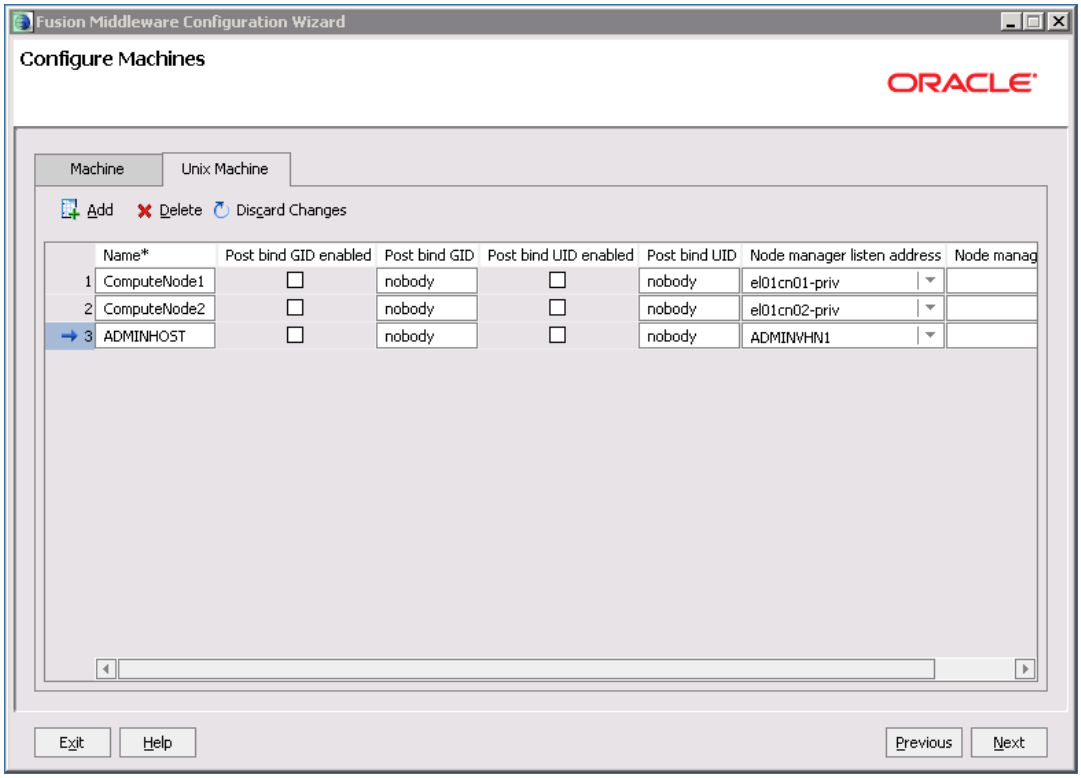

Click **Next**.

The **Assign Servers to Machines** screen is displayed.

14. Perform the following:

- Select **ADMINHOST** in the **Machine** pane, and double-click the **AdminServer**.

  The **AdminServer** is removed from the **Machine** pane and added below **ADMINHOST** in the **Machine** pane (see Figure 5–6).

- Select **ComputeNode1** in the **Machine** pane, and double-click the Managed Servers you created on `ComputeNode1` (See Table 5–2).

  The name of the Managed Server is removed from the **Machine** pane and added below **ComputeNode1** in the **Machine** pane (see Figure 5–6).

- Select **ComputeNode2** in the **Machine** pane, and double-click the Managed Servers you created on `ComputeNode2`.

  The name of the Managed Server is removed from the **Server** pane and added below **ComputeNode2** in the **Machine** pane (see Figure 5–6).

Figure 5–6 Assign Servers to Machines

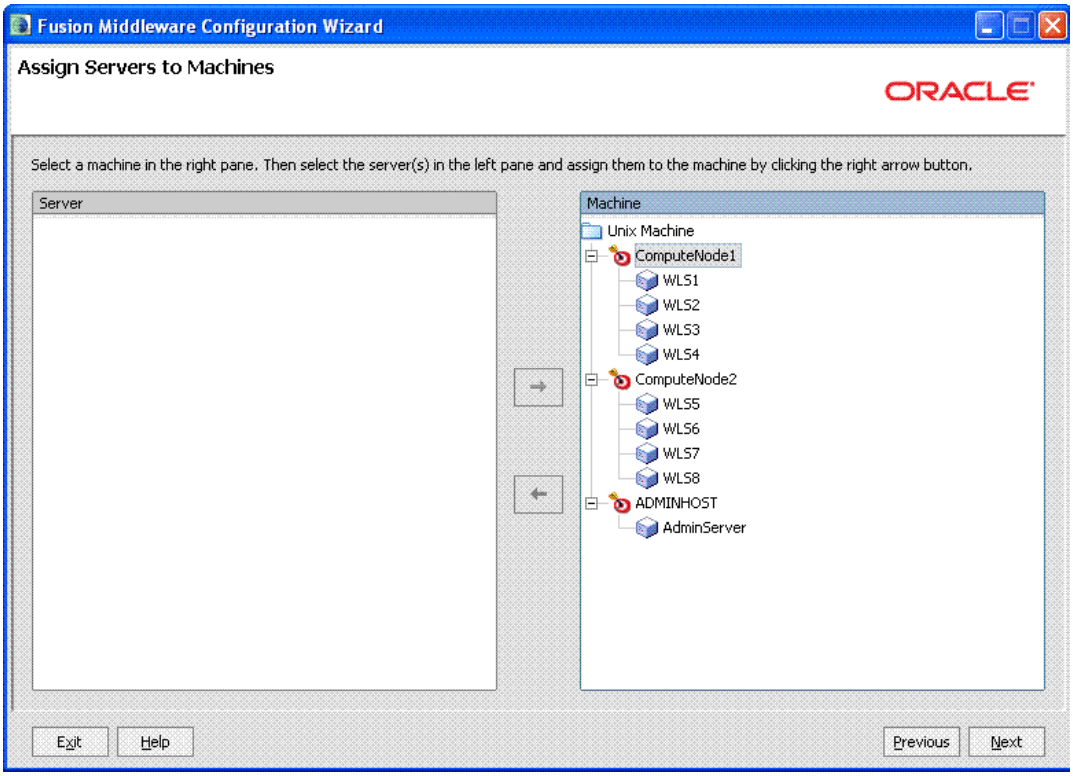

- Click **Next**.

15. The **Configuration Summary** screen is displayed. In this screen, review the summary of configuration you have chosen and click **Create**.
16. In the **Create Domain** screen, click **Done**.

5.5 Creating boot.properties for the Administration Server on ComputeNode1

Create a `boot.properties` file for the Administration Server on `ComputeNode1`. The `boot.properties` file enables the Administration Server to start without prompting you for the administrator username and password. To create a `boot.properties` file, complete the following steps:

1. Create the following directory structure:

   ```
   mkdir -p /u01/Dept_1/domains/el01cn01/base_domain/servers/AdminServer/security
   ```

2. In a text editor, create a file called `boot.properties` in the directory created in the previous step, and enter the following lines in the file:

   ```
   username=<adminuser>
   password=<password>
   ```

Note: When you start the Administration Server, the username and password entries in the file get encrypted. You start the Administration Server in Section 5.6, "Starting the Administration Server on ComputeNode1."

For security reasons, you want to minimize the time the entries in the file are left unencrypted: after you edit the file, you should start the server as soon as possible so that the entries get encrypted.

5.6 Starting the Administration Server on ComputeNode1

The Administration Server is started and stopped using Node Manager. However, the first start of the Administration Server with Node Manager requires configuration changes that are not set for Node Manager by the Configuration Wizard. Therefore, use the start script for the Administration Server for the first start.

Note: ■Open the `setDomainEnv.sh` file (Located in `/u01/Dept_1/domains/el01cn01/base_domain/bin`) in a text editor and change the memory allocation (`WLS_MEM_ARGS_64BIT="-Xms512m -Xmx512m"`) to `1024m` and `3072m`, as shown in the following example:

```
WLS_MEM_ARGS_64BIT="-Xms1024m -Xmx3072m"
```

1. To start the Administration Server on `ComputeNode1`, run the following script on the command-line:

   ```
   ./startWebLogic.sh
   ```

   `startWebLogic.sh` is located in `/u01/Dept_1/domains/el01cn01/base_domain/bin`.

2. In a browser, go to the following URL:

   ```
   http://ADMINVHN1:7001/console
   ```

   Note: Ensure that you have configured the network before performing the domain configuration for the Administration Server and Managed Servers. The network configuration is described in Section 3.3, "Network".

3. Log in to the Administration Server Console using the WebLogic Administration Server user name and password.
4. In the banner toolbar region at the top of the right pane of the Console, click **Preferences**.

   The **Preferences** page is displayed.
5. Click **Shared Preferences**, and then deselect **Follow Configuration changes**.
6. Click **Save**.

5.7 Configuring Java Node Manager

Node Manager is an Oracle WebLogic Server utility that enables you to start, shut down, and restart Administration Server and Managed Server instances. Node manager does the following:

- Starts a Managed Server through the Administration Console.
- Monitors Servers that it has started.
- If a Managed Server shuts down improperly for any reason, then the Automatic restart restarts the Managed Server. Node Manager does not typically kill a server. If you shut down the server using the Administration Console, Node Manager, or WLST, then Node Manager would not automatically restart that server. It starts the server only when the process stops unexpectedly.

Java-based Node Manager runs within a Java Virtual Machine (JVM) process. On UNIX platforms, allowing it to restart automatically when the system is rebooted.

Recommendations

Oracle provides the following recommendations for Node Manager configuration in enterprise deployment topologies:

1. Place the Node Manager configuration and log file in a location different from the default one (which is inside the Middleware Home where Node Manager resides). See Section 5.7.2, "Changing the Location of Node Manager Configuration Files" for further details.
2. Set up a dedicated Node Manager for each of the compute nodes.
3. Every Node Manager instance running on the respective compute nodes must have a dedicated Node Manager home directory.

This section describes how to configure Node Manager for your Oracle Exalogic enterprise deployment implementation. You must complete the following tasks:

- Starting Node Manager to Generating the Properties File
- Changing the Location of Node Manager Configuration Files
- Editing nodemanager.properties File
- Specifying Node Manager Username and Password
- Starting Node Manager
- Controlling and Configuring Node Manager Using WLST

5.7.1 Starting Node Manager to Generating the Properties File

You must start the Node Manager to generate the properties file. Perform these steps to start Node Manager:

1. Start Node Manager:

```
ComputeNode1> cd /u01/FMW_Product1/wlserver_10.3/server/bin
ComputeNode1> ./startNodeManager.sh
```

2. Open `nodemanager.properties` (Located at `u01/FMW_Product1/wlserver_10.3/common/nodemanager` directory), and set the `StartScriptEnabled` property to `true`.

Note: You must use the `StartScriptEnabled` property to avoid class loading failures and other problems.

3. Stop the Node Manager process by running the following commands:

   ```
   ps -eaf |grep NodeManager
   ```

   Example output:

   ```
   user    10597   472  7 10:40 pts/3    00:00:00 java
   weblogic.NodeManager
   ```

   Run the `kill` command to stop the Node Manager process, as in the following example:

   ```
   kill 10597
   ```

4. Repeat these steps on `ComputeNode2`.

5.7.2 Changing the Location of Node Manager Configuration Files

You must create a new directory for Node Manager configuration and log files outside the `MW_HOME` directory, and perform all Node Manager configuration tasks from this directory. Ensure that you do not make any configuration changes to the Node Manager files located in the Oracle WebLogic home directory.

You must complete the following steps:

1. Create the following directories:

   On `ComputeNode1`:

   ```
   mkdir -p /u01/Dept_1/admin/el01cn01/nodemanager
   ```

   On `ComputeNode2`:

   ```
   mkdir -p /u01/Dept_1/admin/el01cn02/nodemanager
   ```

2. Go to the **nodemanager** folder in your `/u01/FMW_Product1/wlserver_10.3/common/nodemanager` directory and copy the `nodemanager.properties` file to the new **nodemanager** folder you created for `ComputeNode1` and `ComputeNode2`:

3. Copy `startNodeManager.sh` (located in the `/u01/FMW_Product1/wlserver_10.3/server/bin` directory) and `nodemanager.domains` files (located in the `/u01/FMW_Product1/wlserver_10.3/common/nodemanager` directory) to the new **nodemanager** folder you created for `ComputeNode1` and `ComputeNode2`.

4. Open `startNodeManager.sh` for `ComputeNode1` and `ComputeNode2` (Located in your new nodemanager folder in `ComputeNode1` and `ComputeNode2`) using a text editor, and make the following change:

   On `ComputeNode1`:

   ```
   NODEMGR_HOME="/u01/Dept_1/admin/el01cn01/nodemanager"
   ```

   On `ComputeNode2`:

   ```
   NODEMGR_HOME="/u01/Dept_1/admin/el01cn02/nodemanager"
   ```

Note: The `startNodeManager.sh` script can now be used to start the Node Manager on `ComputeNode1` and `ComputeNode2`.

5.7.3 Editing nodemanager.properties File

Node Manager properties define a variety of configuration settings for a Java-based Node Manager process. You must specify Node Manager properties on the command line or define them in the `nodemanager.properties` file (Located at `/u01/Dept_1/admin/el01cn01/nodemanager` on `ComputeNode1` and `/u01/Dept_1/admin/el01cn02/nodemanager` on `ComputeNode2`). Values supplied on the command line override the values in `nodemanager.properties`.

Table 5–4 lists the Node Manager properties that you must change for `ComputeNode1`.

Table 5–4 Node Manager Properties

Properties	Value
`SecureListener`	Set the value to "`false`".
`StartScriptEnabled`	Set the value to "`true`", to start a server. For more information, see the section "Configuring Node Manager to Use Start and Stop Scripts" in the *Oracle Fusion Middleware Node Manager Administrator's Guide for Oracle WebLogic Server*.
`StopScriptEnabled`	Set the value to "`true`", to run a server after a server is stopped, killed, or crashed. For more information, see the section "Configuring Node Manager to Use Start and Stop Scripts" in the *Oracle Fusion Middleware Node Manager Administrator's Guide for Oracle WebLogic Server*.
`StopScriptName`	Specify a name for the stop script, for example `stopWebLogic.sh`. For more information, see the section "Configuring Node Manager to Use Start and Stop Scripts" in the *Oracle Fusion Middleware Node Manager Administrator's Guide for Oracle WebLogic Server*.
`DomainsFile`	`/u01/Dept_1/admin/el01cn01/nodemanager/nodemanager.domains`
`ListenAddress`	`192.168.10.1`
`NodeManagerHome`	`/u01/Dept_1/admin/el01cn01/nodemanager`
`LogFile`	`/u01/Dept_1/admin/el01cn01/nodemanager/nodemanager.log`

You must also edit the `nodemanager.properties` file for `ComputeNode2` as described in Table 5–4, and ensure that you enter the following values:

- `NodeManagerHome`: `/u01/Dept_1/admin/el01cn02/nodemanager`
- `ListenAddress=` `192.168.10.2`

Note: This IP address is the `Bond0` IP address of `ComputeNode2`.

- `LogFile=` `/u01/Dept_1/admin/el01cn02/nodemanager/nodemanager.log`

- DomainsFile= /u01/Dept_1/admin/el01cn02/nodemanager/nodemanager.domains

> **Note:** For more information about nodemanager.properties, see the section "Reviewing nodemanager.properties" in the *Oracle Fusion Middleware Node Manager Administrator's Guide for Oracle WebLogic Server*.

5.7.4 Specifying Node Manager Username and Password

Use the Administration Console to update the Node Manager credentials for ComputeNode1 and ComputeNode2:

1. In a browser, go to the following URL:

   http://ADMINVHN1:7001/console

2. Log in as the administrator.
3. Click **Lock and Edit**.
4. In the left pane of the Console, select **base_domain**. This is the domain you have created for the Oracle Exalogic enterprise deployment.

   The **Settings for base_domain** page is displayed.
5. Click **Security** tab, and then **General** tab.
6. Expand the **Advanced** options at the bottom.
7. Enter a new username for Node Manager, or make a note of the existing one and update the Node Manager password.

   > **Note:** You need the Node Manager credentials to connect Node Manager with nmconnect.

8. Click **Save**.
9. Click **Activate Changes**.

5.7.5 Starting Node Manager

To start the Node Manager on ComputeNode1, run startNodeManager.sh (Located at /u01/Dept_1/admin/el01cn01/nodemanager directory for ComputeNode1 and /u01/Dept_1/admin/el01cn02/nodemanager directory for ComputeNode2) command to start Node Manager:

On ComputeNode1:

```
ComputeNode1> cd /u01/Dept_1/admin/el01cn01/nodemanager
ComputeNode1> ./startNodeManager.sh
```

On ComputeNode2:

```
ComputeNode2> cd /u01/Dept_1/admin/el01cn02/nodemanager
ComputeNode2> ./startNodeManager.sh
```

5.7.6 Controlling and Configuring Node Manager Using WLST

The WebLogic Scripting Tool (WLST) is a command-line scripting interface that system administrators and operators use to monitor and manage WebLogic Server instances and domains. You can start, stop, and restart server instances remotely or locally, using WLST as a Node Manager client. In addition, WLST can obtain server status and retrieve the contents of the server output log and Node Manager log. For more information on WLST commands, see "WLST Command and Variable Reference" in *WebLogic Scripting Tool Command Reference*.

Using nmConnect in a Production Environment

WLST can connect to a Node Manager that is running on any machine and start one or more WebLogic Server instances on the machine. A domain's Administration Server does not need to be running for WLST and Node Manager to start a server instance using this technique.

However, by default, the `nmConnect` command cannot be used in a production environment. You must first perform the following procedures to use `nmConnect` in a production environment.

1. Start WLST as follows:

```
ComputeNode1> cd /u01/app/FMW_Product1/Oracle/Middleware/wlserver_
10.3/common/bin
ComputeNode1> ./wlst.sh
```

2. Use the `connect` command to connect WLST to a WebLogic Server instance, as in the following example:

```
wls:/offline> connect('username','password','t3://ADMINVHN1:7001')
```

3. Once you are in the WLST shell, run `nmEnroll` using the following syntax:

```
nmEnroll([domainDir], [nmHome])
```

   For example,

```
nmEnroll('/u01/Dept_1/domains/el01cn01/base_domain', '/u01/Dept_
1/admin/el01cn01/nodemanager')
```

   Running `nmEnroll` ensures that the correct Node Manager user and password token are supplied to each Managed Server. Once these are available for each Managed Server, you can use `nmConnect` in a production environment.

4. Disconnect WLST from the WebLogic Server instance by entering `disconnect()`, and exit by entering `exit()` to exit the WLST shell.

> **Note:** You must run `nmEnroll` for both `ComputeNode1` and `ComputeNode2`.

5.8 Restarting the Administration Server on ComputeNode1

You must stop the Administration Server, and restart it using the Node Manager. Complete the following steps:

1. Stop the Administration Server process by using **CTRl-C** in the shell where it was started, or by process identification and kill in the Operating System. Alternatively, you can stop the Administration Server by using the WebLogic Administration Console.

2. Start WLST and connect to Node Manager with **nmconnect** and the credentials set in Section 5.7.4, "Specifying Node Manager Username and Password" and start the Administration Server using **nmstart**.

```
ComputeNode1> cd /u01/FMW_Product1/wlserver_10.3/common/bin
ComputeNode1> ./wlst.sh
```

Once you are in the WLST shell:

```
wls:/offline>nmConnect('Admin_User','Admin_Password',
'el01cn01-priv','5556','base_domain','/u01/Dept_1/domains/el01cn01/base_
domain','ssl')

wls:/nm/base_domain> nmStart('AdminServer')
```

Note: The username and password are used only to authenticate connections between Node Manager and clients. They are independent of the server admin ID and password. This is the user you specified in Specifying Node Manager Username and Password.

5.9 Validating the Administration Server

Perform these steps to ensure that the Administration Server is properly configured:

1. In a browser, go to `http://ADMINVHN1:7001/console`.
2. Log in as the administrator.
3. In the left pane of the Console, expand **Environment**, and then **Servers**.

   The **Summary of Servers** page is displayed.
4. In the Summary of Servers, ensure that the Managed Servers created for `ComputeNode1` and `ComputeNode2` are listed (Figure 5–7).

Figure 5–7 Summary of Servers

New | Clone | Delete — Showing 1 to 9 of 9 Previous | Next

Name	Cluster	Machine	State	Health	Listen Port
AdminServer(admin)		ADMINHOST	RUNNING	OK	7001
WLS1	Dept1_Cluster1	ComputeNode1	SHUTDOWN		7003
WLS2	Dept1_Cluster1	ComputeNode1	SHUTDOWN		7003
WLS3	Dept1_Cluster1	ComputeNode1	SHUTDOWN		7003
WLS4	Dept1_Cluster1	ComputeNode1	SHUTDOWN		7003
WLS5	Dept1_Cluster1	ComputeNode2	SHUTDOWN		7003
WLS6	Dept1_Cluster1	ComputeNode2	SHUTDOWN		7003
WLS7	Dept1_Cluster1	ComputeNode2	SHUTDOWN		7003
WLS8	Dept1_Cluster1	ComputeNode2	SHUTDOWN		7003

New | Clone | Delete — Showing 1 to 9 of 9 Previous | Next

5.10 Optional: Creating Oracle HTTP Server Instances in the Exalogic Environment

You should complete this task if you wish to use Oracle HTTP Server to load balance traffic on the private InfiniBand network (`BOND0/IPMP0`). Before configuring Oracle

HTTP Server to handle this load balancing, you should have installed Oracle HTTP Server in the Exalogic environment. This installation is different than the Oracle HTTP Server installation running outside of Exalogic. Be sure to use the "*internal*" Oracle HTTP Server for load balancing the application traffic on the private network only. For handling external requests in the public DMZ, you should continue to use the Oracle HTTP Server running outside of Exalogic.

To set up Oracle HTTP Server instances in the Exalogic environment, complete the following steps:

1. Ensure that Oracle HTTP Server is installed, as described in Section 4.3, "Installing Oracle HTTP Server".
2. On an Exalogic compute node (for example, `ComputeNode1`), create Oracle HTTP Server instances as follows:

   Note: In the physical topology, you may create two Oracle HTTP Server instances on an Exalogic compute node.

   a. From the Oracle HTTP Server Home directory (`/u01/app/FMW_Product1/Oracle/Middleware/Oracle_WT1`), run the following command to start the Configuration Wizard:

      ```
      # cd /bin/config.sh
      ```

      The **Welcome** screen is displayed.

   b. Click **Next**.

      The **Configure Components** Screen is displayed.

   c. Select **Oracle HTTP Server** and **Select Associate Selected Components with WebLogic Domain**.

      Click **Next**.

      The **Specify WebLogic Domain** Screen is displayed.

   d. Specify the following WebLogic Domain credentials:

      - **Domain Host Name**: `base_domain`
      - **Domain Port No**: `7001`
      - **User Name**: Specify the user name. The default user name is `weblogic`.
      - **Password**: Specify the Administration Server password.

      Click **Next**.

      The **Specify Component Details** Screen is displayed.

   e. Specify the following component details:

      - **Instance Home Location**: `/u01/app/FMW_Product1/Oracle/Middleware/Oracle_WT1/instances/`
      - **Instance Name**: `asinst_1`
      - **OHS Component Name**: OHS1

      Click **Next**.

      The **Configure Ports Screen** is displayed.

   f. Select **Auto Port Configuration** to automatically assign the default port.

Click **Next**.

The **Specify Security Updates Screen** is displayed.

g. Enter your E-mail address if you want to receive the latest product information and security updates. If you have a My Oracle account and wish to receive updates via this mechanism, select **I wish to receive security updates via My Oracle Support**, then enter your account password.

If you do not wish to register for security updates, leave all the fields on this screen blank. You will be prompted to confirm your selection. Click **Yes** to confirm that you do not want to register for any security updates.

Click **Next**.

The **Installation Summary Screen** is displayed.

h. Review the information on this screen. The operations summarized on this page will be performed when you click **Configure**.

The **Configuration Progress Screen** is displayed.

i. Click **Next**.

The **Installation Complete Screen** is displayed.

You can also save this summary information to a file for future reference by clicking **Save**. You will be prompted to specify a name and location for your summary file.

j. Click **Finish** to dismiss the screen.

An Oracle HTTP Server instance named `OHS1` is created under the `/u01/app/FMW_Product1/Oracle/Middleware/Oracle_WT1/instances` directory.

3. Edit the configuration files for the `OHS1` instance as follows:

a. Open the `/u01/app/FMW_Product1/Oracle/Middleware/Oracle_WT1/instances/OHS1/config/OHS/httpd.conf` file in a text editor. search for the `Listen` directive, which is usually commented out. Uncomment the `Listen` entry and specify a `BOND0/IPMP0` private IP address, such as `10.0.0.18`.

b. Save the `httpd.conf` file and close.

c. Open the `/u01/app/FMW_Product1/Oracle/Middleware/Oracle_WT1/instances/OHS1/config/OPMN/opmn/opmn.xml` file in a text editor. Add the following property in the opmn.xml file:

```
<ipaddr  remote="ip" request="ip"/>
```

d. Specify the `BOND0/IPMP0` private IP address, such as `10.0.0.19`, for the `request` attribute as follows:

```
<ipaddr  remote="ip" request="10.0.0.19"/>
```

e. Save the `opmn.xml` file and close.

Note: You must repeat these steps to create Oracle HTTP Server instances to run on the remaining compute nodes, such as `ComputeNode2`, in the example configuration scenario.

For load balancing application traffic in a round-robin fashion on the `BOND0/IPMP0` private network, you must configure Oracle HTTP Server, as described in Section 6.4, "Optional: Configuring Oracle HTTP Server for Load Balancing on the Private InfiniBand Network".

5.11 Propagating Domain Configuration from ComputeNode1 to ComputeNode2 Using pack and unpack Utilities

You have created the domain (`base_domain`) on `ComputeNode1`. You must propagate the domain configuration to `ComputeNode2` as follows:

1. Run the `pack` command on `ComputeNode1` to create a template pack using the following commands:

```
ComputeNode1> cd /u01/app/FMW_Product1/Oracle/Middleware/wlserver_
10.3/common/bin

ComputeNode1> ./pack.sh -managed=true -domain=/u01/Dept_
1/domains/el01cn01/base_domain -template=basedomaintemplate.jar -template_
name=basedomain_template
```

2. Run the `unpack` command on `ComputeNode2` to unpack the template.

```
ComputeNode2> cd /u01/app/FMW_Product1/Oracle/Middleware/wlserver_
10.3/common/bin

ComputeNode2> ./unpack.sh -domain=/u01/Dept_1/domains/el01cn02/base_domain
-template=basedomaintemplate.jar
```

5.12 Configuring Network Channels for HTTP and T3 Clients via EoIB

If your HTTP clients and T3 clients use the 10 Gb Ethernet network, you must create additional network channels for the Administration Server and Managed Servers on `ComputeNode1` and `ComputeNode2`. For more information, see Section 3.3.3.2, "Determining Network Interface and Channel Requirements for a WebLogic Managed Server and the Administration Server".

Complete the following steps:

- Section 5.12.1, "Using BOND1 Floating IP Addresses for EoIB Communication"
- Section 5.12.2, "Configuring the Administration Server Network Channel"
- Section 5.12.3, "Configuring Network Channels for Managed Servers on ComputeNode1 and ComputeNode2"

5.12.1 Using BOND1 Floating IP Addresses for EoIB Communication

For Ethernet over InfiniBand (EoIB) communication, you must log in as a `root` user and assign floating IP addresses to the Administration Server and to Managed Servers. These addresses are associated with the `BOND1` interface. In this guide, an IP subnet (`10.1.0.0`) is used as an example only. For more information, see Section 3.3.3.4, "IP Addresses for WebLogic Clusters When HTTP or T3 Traffic Is Via Ethernet over InfiniBand (EoIB)".

5.12.2 Configuring the Administration Server Network Channel

You must create two network channels for the Administration Server on `ComputeNode1`. The network channels are necessary for routing HTTP traffic and T3 traffic (TCP-based protocol used by WebLogic Server) coming in from external data center via Ethernet over InfiniBand (EoIB). You must create the following network channels:

- HTTP Client Channel
- T3 Client Channel

5.12.2.1 HTTP Client Channel

To create the HTTP network channel for the Administration Server, complete the following steps:

1. In a browser, go to `http://ADMINVHN1:7001/console`.
2. Log in as the administrator.
3. If you have not already done so, click **Lock & Edit** in the Change Center.
4. In the left pane of the Console, expand **Environment**, and then **Servers**.

   The **Summary of Servers** page is displayed.
5. In the Servers table, click **AdminServer(admin)**.

   The **Settings for AdminServer** page is displayed.
6. Select **Protocols**, and then **Channels**.
7. Click **New**.
8. Enter **AdminHTTPClient** as the name of the new network channel and select **http** as the protocol, then click **Next**.
9. Enter the following information in the **Network Channel Addressing** page:
   - Listen address: **10.1.0.17**

     Note: This address is the floating IP assigned to the Administration Server using the `BOND1` interface.
   - Listen port: **7001**
10. Click **Next**, and select **Enabled** on the **Network Channel Properties** page.
11. Click **Finish**.
12. To activate these changes, in the Change Center of the Administration Console, click **Activate Changes**.

5.12.2.2 T3 Client Channel

To create the t3 network channel for the Administration Server, complete the following:

1. In a browser, go to `http://ADMINVHN1:7001/console`.
2. Log in as the administrator.
3. If you have not already done so, click **Lock & Edit** in the Change Center.
4. In the left pane of the Console, expand **Environment**, and then **Servers**.

The **Summary of Servers** page is displayed.

5. In the Servers table, click **AdminServer(admin)**.

   The **Settings for AdminServer** page is displayed.

6. Select **Protocols**, and then **Channels**.
7. Click **New**.
8. Enter **AdminT3** as the name of the new network channel and select **t3** as the protocol, then click **Next**.
9. Enter the following information in the **Network Channel Addressing** page:
   - Listen address: **10.1.0.17**

     Note: This address is the floating IP assigned to the Administration Server using the `BOND1` interface.

   - Listen port: **7001**
10. Click **Next**, and select the following in the **Network Channel Properties** page:
    - **Enabled**
    - **HTTP Enabled for This Protocol**
11. Click **Finish**.
12. To activate these changes, in the Change Center of the Administration Console, click **Activate Changes**.

5.12.3 Configuring Network Channels for Managed Servers on ComputeNode1 and ComputeNode2

For Managed Servers, you must create the following network channels using Ethernet over InfiniBand (`EoIB`):

- HTTP Client Channel
- T3 Client Channel

5.12.3.1 HTTP Client Channel

To create a HTTP network channel for a Managed Server such as WLS1, complete the following steps:

1. In a browser, go to `http://ADMINVHN1:7001/console`.
2. Log in as the administrator.
3. If you have not already done so, click **Lock & Edit** in the Change Center.
4. In the left pane of the Console, expand **Environment**, and then **Servers**.

   The **Summary of Servers** page is displayed.

5. In the Servers table, click **WLS1**.

   The **Settings for WLS1** page is displayed.

6. Select **Protocols**, and then **Channels**.
7. Click **New**.

8. Enter **HTTPClient** as the name of the new network channel and select **http** as the protocol, then click **Next**.
9. Enter the following information in the **Network Channel Addressing** page:
   - Listen address: **10.1.0.1**

     Note: This address is the floating IP assigned to a Managed Server using the `BOND1` interface.

   - **Listen port**: `7003`
   - **External Listen Address**: `exalogic.mycompany.com`

     Note: This is the IP address or DNS name to access application on the server.

   - **External Listen Port**: `80`
10. Click **Next**, and select the following in the **Network Channel Properties** page:
    - **Enabled**
    - **HTTP Enabled for This Protocol**
11. Click **Finish**.
12. To activate these changes, in the Change Center of the Administration Console, click **Activate Changes**.

You must repeat the above steps to create a network channel each for the remaining Managed Servers on `ComputeNode1` and `ComputeNode2` and enter the required properties as described in Table 5–5.

Note: In this example, IP addresses are used as listen addresses. However, you can specify host names if they resolve to their corresponding floating IP addresses.

Table 5–5 Network Channels Properties

Managed Server	Name	Protocol	Listen Address	Listen Port	External Listen Address	External Listen Port
WLS2	HTTPClient	http	10.1.0.2	7003	exalogic.mycompany.com	80
WLS3	HTTPClient	http	10.1.0.3	7003	exalogic.mycompany.com	80
WLS4	HTTPClient	http	10.1.0.4	7003	exalogic.mycompany.com	80
WLS5	HTTPClient	http	10.1.0.5	7003	exalogic.mycompany.com	80
WLS6	HTTPClient	http	10.1.0.6	7003	exalogic.mycompany.com	80
WLS7	HTTPClient	http	10.1.0.7	7003	exalogic.mycompany.com	80
WLS8	HTTPClient	http	10.1.0.8	7003	exalogic.mycompany.com	80

5.12.3.2 T3 Client Channel

To create the t3 network channel for a Managed Server server such as `WLS1`, complete the following steps:

1. In a browser, go to `http://ADMINVHN1:7001/console`.
2. Log in as the administrator.
3. If you have not already done so, click **Lock & Edit** in the Change Center.
4. In the left pane of the Console, expand **Environment**, and then **Servers**.

   The **Summary of Servers** page is displayed.
5. In the Servers table, click **AdminServer(admin)**.

   The **Settings for AdminServer** page is displayed.
6. Select **Protocols**, and then **Channels**.
7. Click **New**.
8. Enter **T3ClientChannel** as the name of the new network channel and select **t3** as the protocol, then click **Next**.
9. Enter the following information in the **Network Channel Addressing** page:
   - Listen address: **10.1.0.1**

     Note: This address is the floating IP assigned to a Managed Server or to the Administration Server using the `BOND1` interface.
   - Listen port: **7005**
10. Click **Next**, and select the following in the **Network Channel Properties** page:
    - **Enabled**
    - If you want to allow both HTTP and T3 traffic on `T3ClientChannel1`, then select the following option:

      HTTP Enabled for This Protocol
11. Click **Finish**.
12. To activate these changes, in the Change Center of the Administration Console, click **Activate Changes**.

You must repeat the above steps to create a network channel each for the remaining Managed Servers and enter the required properties as described in Table 5–6.

Note: In this example, IP addresses are used as listen addresses. However, you can specify host names if they resolve to their corresponding floating IP addresses.

Table 5–6 Network Channels Properties

Managed Server	Name	Protocol	Listen Address	Listen Port	External Listen Address	External Listen Port
WLS2	T3ClientChannel	t3	10.1.0.2	7005	exalogic.mycompany.com	80
WLS3	T3ClientChannel	t3	10.1.0.3	7005	exalogic.mycompany.com	80
WLS4	T3ClientChannel	t3	10.1.0.4	7005	exalogic.mycompany.com	80
WLS5	T3ClientChannel	t3	10.1.0.5	7005	exalogic.mycompany.com	80
WLS6	T3ClientChannel	t3	10.1.0.6	7005	exalogic.mycompany.com	80
WLS7	T3ClientChannel	t3	10.1.0.7	7005	exalogic.mycompany.com	80
WLS8	T3ClientChannel	t3	10.1.0.8	7005	exalogic.mycompany.com	80

5.13 Configuring Oracle Coherence

This section describes how to configure Oracle Coherence for Oracle Exalogic enterprise deployment. You must complete the following tasks:

- Configuring Socket Buffer Sizes
- Creating Coherence Clusters on ComputeNode1 and ComputeNode2
- Deploying the Coherence Shared Library Files
- Create the Counter Web Application
- Deploy the Application
- Creating Coherence Servers on ComputeNode1 and ComputeNode2
- Configuring Startup Arguments for Coherence servers
- Starting Coherence Servers on ComputeNode1 and ComputeNode2
- Verify the Example

5.13.1 Configuring Socket Buffer Sizes

To help minimization of packet loss, the operating system socket buffers should be large enough to handle the incoming network traffic while your Java application is paused during garbage collection. By default Coherence will attempt to allocate a socket buffer of 2MB. If your operating system is not configured to allow for large buffers Coherence will use smaller buffers. Most versions of UNIX have a very low default buffer limit, which should be increased to at least 2MB.

If the operating system fails to allocate the full size buffer, the following error message is displayed:

Example 5–1 Message Indicating OS Failed to Allocate the Full Buffer Size

```
UnicastUdpSocket failed to set receive buffer size to 1428 packets (2096304
bytes); actual size is 89 packets (131071 bytes). Consult your OS documentation
```

```
regarding increasing the maximum socket buffer size. Proceeding with the actual
value may cause sub-optimal performance.
```

Although it is safe to operate with the smaller buffers, it is recommended that you configure your operating system to allow for larger buffers.

On Oracle Linux, execute (as root):

```
sysctl -w net.core.rmem_max=4192608
sysctl -w net.core.wmem_max=4192608
```

You may configure Coherence to request alternate sized buffers for packet publishers and unicast listeners by using the `coherence/cluster-config/packet-publisher/packet-buffer/maximum-packets` and `coherence/cluster-config/unicast-listener/packet-buffer/maximum-packets` elements. For more information, see "packet-buffer" in the *Oracle Coherence Developer's Guide*.

On Oracle Solaris, execute (as root):

```
ndd -set /dev/udp udp_max_buf 2096304
```

5.13.2 Creating Coherence Clusters on ComputeNode1 and ComputeNode2

You must enable applications to share data management and caching services among Managed Server instances and clusters hosting the applications that require access to them.

In the example configuration scenario discussed in this guide, one Oracle Coherence Cluster (`CoherenceCluster1`) and one Oracle WebLogic Cluster (`Dept1_Cluster1`) are used. You should determine the number of Coherence clusters, based on your application deployment requirements. If you decide to use multiple Coherence clusters, you may want to configure a well-known address (WKA) for the clusters, see "Configure well known addresses" in the *Oracle Fusion Middleware Oracle WebLogic Server Administration Console Online Help*.

This example procedure uses default values only.

> **Note:** The configuration example used in this guide uses two compute nodes for Coherence clusters as a canonical example only. Oracle recommends that you determine the number of physical machines (compute nodes) based on your specific application deployment requirements. For more information, see the "Production Checklist" in the *Oracle Coherence Developer's Guide* for recommendations about deploying Oracle Coherence in a production environment.

To create a Coherence cluster across `Dept1_Cluster1`:

1. Log in to the Oracle WebLogic Administration Console, using the following URL in a browser:

   ```
   http://ADMINVHN1:7001/console
   ```

2. If you have not already done so, click **Lock & Edit** in the Change Center.
3. In the left pane of the Console, expand **Environment** and select **Coherence Clusters**.

The **Summary of Coherence Clusters** page is displayed.

4. Click **New**.

   The **Create Coherence Cluster Configuration** page is displayed.

5. Enter **CoherenceCluster1** as the name of the Coherence cluster configuration and click **Next**.

   The **Coherence Cluster Addressing** page is displayed.

6. Accept the default values, and click **Next**.

   The **Coherence Cluster Targets** page is displayed.

7. Select **Dept1_Cluster1**, and then **All servers in the cluster** (Figure 5–8) to deploy this Coherence cluster configuration.

Figure 5–8 Coherence Cluster Targets

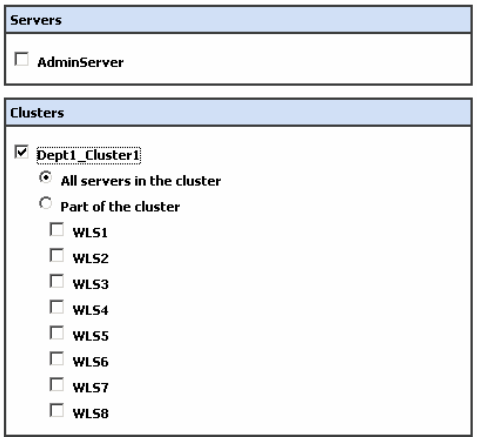

8. Click **Finish**.

The **Summary of Coherence Clusters** page is displayed, displaying the Coherence Cluster you created for `Dept1_Cluster1` (Figure 5–9).

Figure 5–9 Summary of Coherence Clusters

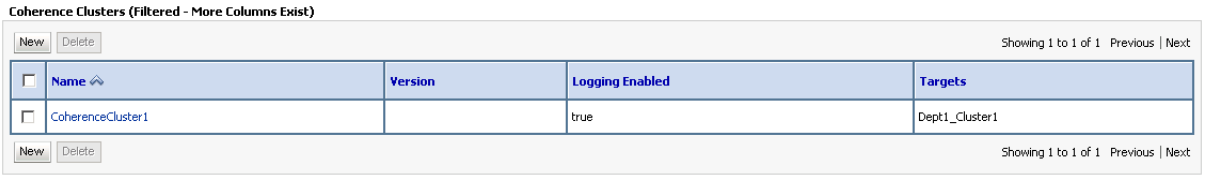

5.13.3 Deploying the Coherence Shared Library Files

You must deploy the following shared library files:

- coherence-web-spi.war
- active-cache.jar
- coherence.jar

coherence-web-spi.war

Oracle Coherence provides a deployable shared library named `coherence-web-spi.war` that contains a native plug-in to WebLogic Server's HTTP Session Management interface. You must deploy this file as follows:

1. Log in to the Oracle WebLogic Administration Console.
2. If you have not already done so, click **Lock & Edit** in the Change Center.
3. In the left pane of the Console, click **Deployments**.

   The **Summary of Deployments** page is displayed.
4. Click **Install**. The **Install Application Assistant** screen is displayed.
5. Use the **Install Application Assistant** to deploy `coherence-web-spi.war` as a library to `Dept1_Cluster1`. Do the following:
   a. Locate and select the `coherence-web-spi.war` file. It resides in the `/u01/app/FMW_Product1/Oracle/Middleware/coherence_3.6/lib` directory.
   b. Click **Next**.

      The **Choose targeting style** page is displayed.
   c. Ensure that **Install this deployment as a library** is selected.
   d. Click **Next**.

      The **Select deployment targets** page is displayed.
   e. Select **Dept1_Cluster1** as the deployment target.
   f. Click **Next**.
   g. In **Optional Settings** page, select **Copy this application onto every target for me** option in the Source accessibility section.
   h. Click **Finish**.

      The **Summary of Deployments** page is displayed, displaying the `coherence-web-spi.war` file you have installed.
   i. Click **Activate Changes**.

Note: You must extract the contents of the coherence-web-spi.war file to a directory using the following command:

```
jar -xvf coherence-web-spi.war
```

In this example procedure, create a directory named `coherenceweb` in the following directory:

/u01/app/FMW_Product1/Oracle/Middleware/coherence_3.6/lib

---

active-cache.jar

To employ ActiveCache functionality in your applications, you must also deploy the `active-cache-1.0.jar` file to `Dept1_Cluster1`. To do so, complete the following steps:

1. Log in to the Oracle WebLogic Administration Console.
2. If you have not already done so, click **Lock & Edit** in the Change Center.
3. In the left pane of the Console, click **Deployments**.

The **Summary of Deployments** page is displayed.

4. Click **Install**.

   The **Install Application Assistant** screen is displayed.

5. Use the **Install Application Assistant** to deploy `active-cache-1.0.jar` as a library to `Dept1_Cluster1`. Do the following:

   a. Locate and select the `active-cache-1.0.jar` file. It resides in the `/u01/app/FMW_Product1/Oracle/Middleware/wlserver_10.3/common/deployable-libraries` directory.

   b. Click **Next**.

      The **Choose targeting style** page is displayed.

   c. Ensure that **Install this deployment as a library** is selected.

   d. Click **Next**.

      The **Select deployment targets** page is displayed.

      Note: You may get the following warning message, **Issues were encountered while parsing this deployment to determine module type. Assuming this is a library deployment**. This message will appear if the **Choose targeting style** page is skipped automatically by the console. You must ignore this message and proceed.

   e. Select **Dept1_Cluster1** as the deployment target.

   f. Click **Next**.

      The **Optional Settings** page is displayed.

   g. Select **I will make the deployment accessible from the following location** option in the Source accessibility section.

   h. Click **Finish**.

      The **Summary of Deployments** page is displayed, displaying the `active-cache-1.0.jar` file you have installed.

   i. Click **Activate Changes**.

coherence.jar

You must deploy the `coherence.jar` file to `Dept1_Cluster1`. To do so, complete the following steps:

1. Log in to the Oracle WebLogic Administration Console.
2. If you have not already done so, click **Lock & Edit** in the Change Center.
3. In the left pane of the Console, click **Deployments**.

   The **Summary of Deployments** page is displayed.

4. Click **Install**.

   The **Install Application Assistant** screen is displayed.

5. Use the **Install Application Assistant** to deploy `coherence.jar` as a library to `Dept1_Cluster1`. Do the following:

a. Locate and select the `coherence.jar` file. It resides in the `/u01/app/FMW_Product1/Oracle/Middleware/coherence_3.6/lib` directory.

b. Click **Next**.

   The **Choose targeting style** page is displayed.

c. Ensure that **Install this deployment as a library** is selected.

d. Click **Next**.

   The **Select deployment targets** page is displayed.

Note: You may get the following warning message, **Issues were encountered while parsing this deployment to determine module type. Assuming this is a library deployment**. This message will appear if the **Choose targeting style** page is skipped automatically by the console. You must ignore this message and proceed.

---

e. Select **Dept1_Cluster1** as the deployment target.

f. Click **Next**.

   The **Optional Settings** page is displayed.

g. Select **Copy this application onto every target for me** option in the Source accessibility section.

h. Click **Finish**.

   The **Summary of Deployments** page is displayed, displaying the `coherence.jar` file you have installed.

i. Click **Activate Changes**.

5.13.4 Create the Counter Web Application

The Counter Web application is a simple counter implemented as a JSP. The counter is stored as an HTTP session attribute and increments each time the page is accessed.

To create the Counter Web application:

1. Create a standard Web application directory as follows:

```
/
/WEB-INF
```

2. Copy the following code to a text file and save it as a file named `web.xml` in the `/WEB-INF` directory.

```
<?xml version = '1.0' encoding = 'windows-1252'?>
<web-app xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
 xsi:schemaLocation="http://java.sun.com/xml/ns/j2ee
http://java.sun.com/xml/ns/j2ee/web-app_2_4.xsd"
 xmlns="http://java.sun.com/xml/ns/j2ee" version="2.5">
   <description>Empty web.xml file for Web Application</description>
</web-app>
```

3. Create a `weblogic.xml` file in the `/WEB-INF` directory.
   - Add a library reference for the `coherence-web-spi.war` file.
   - Reference the Coherence Cluster in a `coherence-cluster-ref` stanza.

Example 5–2 illustrates a sample `weblogic.xml` file.

Example 5–2 Sample weblogic.xml File

```
<weblogic-web-app xmlns="http://xmlns.oracle.com/weblogic/weblogic-web-app"
xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:schemaLocation="http://xmlns.oracle.com/weblogic/weblogic-web-app
http://www.oracle.com/technology/weblogic/weblogic-web-app/1.1/weblogic-web-app.xs
d">
      <library-ref>
           <library-name>coherence-web-spi</library-name>
      </library-ref>
     <coherence-cluster-ref>
        <coherence-cluster-name>CoherenceCluster1</coherence-cluster-name>
     </coherence-cluster-ref>
</weblogic-web-app>
```

4. Bundle the `coherence.jar` file with the application by copying `coherence.jar` from the `/u01/app/FMW_Product1/wlserver_1034/coherence_3.6/lib` directory to the `WEB-INF/lib` directory.
5. Copy the following code for the counter JSP to a text file and save the file as `counter.jsp` in the root of the Web application directory.

```
<html>
    <body>

<h3>
      Counter :
      <%
         Integer counter = new Integer(1);
         HttpSession httpsession = request.getSession(true);
         if (httpsession.isNew()) {
                httpsession.setAttribute("count", counter);
                out.println(counter);
         } else {
                int count = ((Integer)
httpsession.getAttribute("count")).intValue();
                httpsession.setAttribute("count", new Integer(++count));
                out.println(count);
         }
      %>
      </h3>

    </body>
</html>
```

6. Create a `manifest.mf` file in the `META-INF` directory. Add references to the `active-cache` JAR file. Example 5–3 illustrates a sample `manifest.mf` file.

Example 5–3 Sample manifest.mf File

```
Extension-List: active-cache
active-cache-Extension-Name: active-cache
active-cache-Specification-Version: 1.0
active-cache-Implementation-Version: 1.0
```

7. The Web application directory appears as follows:

```
/
/counter.jsp
```

```
/META-INF/manifest.mf
/WEB-INF/web.xml
/WEB-INF/weblogic.xml
/WEB-INF/lib/coherence.jar
```

8. ZIP or JAR the Web application directory and save the file as `counter.war`.

   Tip: To create an executable JAR file, run the following command from the application's root directory:

   ```
   jar cmf META-INF/manifest.mf ../counter.war
   ```

   This command is an example only.

5.13.5 Deploy the Application

To deploy the `counter.war` application:

1. Open the **Summary of Deployments** page by clicking **Deployments** in the **Domain Structure** menu in the Oracle WebLogic Server Administration Console.
2. Click **Install**. The **Install Application Assistant** wizard opens.
3. Use the **Install Application Assistant** to deploy the `counter.war` file to `Dept1_Cluster1` (all servers in the cluster). In the **Optional Settings** page, select the **Copy this application onto every target for me** option in the **Source accessibility** section.

5.13.6 Creating Coherence Servers on ComputeNode1 and ComputeNode2

You must create four Coherence servers across `ComputeNode1` and `ComputeNode2`, as illustrated in Figure 5–1:

1. If you have not already done so, click **Lock & Edit** in the Change Center of the Administration Console.
2. In the left pane of the Console, expand **Environment** and select **Coherence Servers**.

   The **Summary of Coherence Servers** page is displayed.
3. Click **New**.

   The **Create Coherence Server Configuration** page is displayed.
4. On the Coherence Server Properties page, enter the following:
   - Enter **Coh1** as the name of the server in the **Name** field.

     The server name is not used as part of the URL for applications that are deployed on the server. It is for your identification purposes only. The server name displays in the Administration Console, and if you use WebLogic Server command-line utilities or APIs, you use this name to identify the server.
   - Select **ComputeNode1** as the name of the **Machine** from the drop-down list.

     In order to use Node Manager to start this Coherence server, you must assign the server to a machine.
   - Select **CoherenceCluster1** as the name of the **Cluster** from the drop-down list.
   - **Unicast Listen Address**: Enter `192.168.10.1`

Note: This IP address is the `BOND0` IP of the compute node on which the Oracle Coherence Server is configured to run. Alternatively, you can use the host name of the compute node, such as `el01cn01`, if the host name resolves to the IP address correctly.

For more information, see Section 3.3.3, "Enterprise Deployment Network Configuration".

- Unicast Listen Port: Enter **8888**

Note: Ensure that the Coherence cluster Unicasting Listen Port number specified is different from Coherence Server Unicasting Listen Port number.

- Ensure that **Unicast Port Auto Adjust** is selected.
- Click **Finish**.

5. Repeat the above steps to create seven Coherence Servers on `ComputeNode1` and `ComputeNode2` with the following properties described in Table 5–7:

Note: In this example, IP addresses are used as listen addresses. However, you can specify host names if they resolve to their corresponding floating IP addresses.

Table 5–7 Coherence Server Properties

Name	Machine	Clusters	Unicast Listen Address	Unicast Listen Port
Coh2	ComputeNode1	CoherenceCluster1	192.168.10.1	8890
Coh3	ComputeNode2	CoherenceCluster1	192.168.10.2	8888
Coh4	ComputeNode2	CoherenceCluster1	192.168.10.2	8890

The **Summary of Coherence Servers** page is displayed, displaying the Coherence Servers you have created for `ComputeNode1` and `ComputeNode2` (Figure 5–10).

Figure 5–10 Summary of Coherence Servers

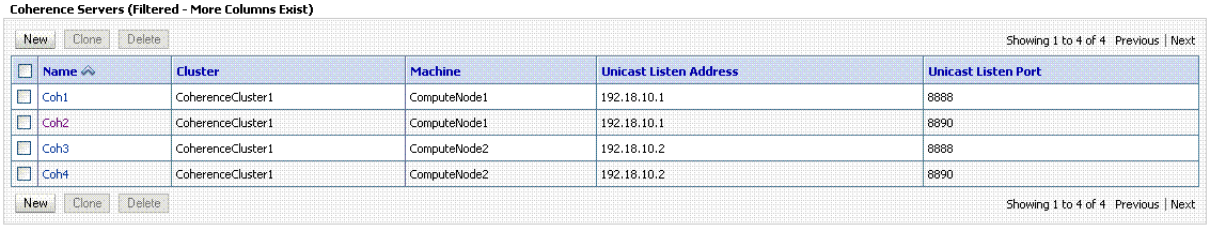

To activate these changes, in the Change Center of the Administration Console, click **Activate Changes**.

5.13.7 Configuring Startup Arguments for Coherence servers

You must configure the startup settings for the Coherence Servers created on `ComputeNode1` and `ComputeNode2`. This setup enables the Node Manager to start a Coherence server. To configure startup options for a Coherence server such as `Coh1`, complete the following steps:

1. If you have not already done so, in the Change Center of the Administration Console, click **Lock & Edit**.
2. In the left pane of the Console, expand **Environment** and select **Coherence Servers**.
3. In the Coherence Servers table, select **Coh1**.
4. Select **Configuration** tab and then select **Server Start** tab.
5. Specify the following startup settings:
   - **Java Home** - Specify the Java home directory (path on the machine running Node Manager) to use when starting this server. Specify the parent directory of the JDK's bin directory. For example: `/u01/app/FMW_Product1/Oracle/Middleware/jrockit_160_20_D1.1.0-15`.
   - **Java Vendor** - Specify the Java vendor value to use when starting this server. For example: Oracle.
   - **BEA Home** - Specify the directory on the Node Manager machine under which all of Oracle's WebLogic products were installed. For example, `/u01/app/FMW_Product1/Oracle/Middleware`.
   - **Root Directory** - Specify the directory that this server uses as its root directory. This directory must be on the computer that hosts the Node Manager. If you do not specify a Root Directory value, the domain directory is used by default. For example: `/u01/Dept_1/domains/el01cn01/base_domain`.
   - **Class Path** - Specify the classpath (path on the machine running Node Manager) to use when starting this server. For example:

     Note: You must run `setWLSEnv.sh` (Located at `/u01/app/FMW_Product1/Oracle/Middleware/wlserver_10.3/server/bin`) and copy its classpath properties to the Coherence Server classpath.

     For example:

     ```
     /u01/app/FMW_Product1/Oracle/Middleware/coherence_
     3.6/lib/coherence.jar:/u01/app/FMW_Product1/Oracle/Middleware/coherence_
     3.6/lib/coherenceweb/WEB_INF/lib/coherence-web.jar:/u01/app/FMW_
     Product1/Oracle/Middleware/modules/features/weblogic.server.modules.coheren
     ce.server_10.3.4.0.jar
     ```

     You must append this classpath to the **Class Path** field.
   - **Arguments** - Specify the arguments to use when starting this server. For example:

     Oracle Linux:

     ```
     -Xms1024m -Xmx3072m
     -Dtangosol.coherence.cacheconfig=session-cache-config.xml
     -Dtangosol.coherence.session.localstorage=true
     -Dtangosol.coherence.cluster=CoherenceCluster1
     ```

Oracle Solaris:

```
-Xms1024m -Xmx3072m
-Dtangosol.coherence.cacheconfig=session-cache-config.xml
-Dtangosol.coherence.session.localstorage=true
-Dtangosol.coherence.cluster=CoherenceCluster1
-Djava.net.preferIPv4Stack=true
```

6. Repeat these steps for the remaining Coherence Servers.

5.13.8 Starting Coherence Servers on ComputeNode1 and ComputeNode2

To start the Coherence Servers on `ComputeNode1` and `ComputeNode2`, complete the following steps:

1. In the WebLogic Administration Server Console, from the **Domain Structure** menu, expand **Environment** and select **Coherence Servers**.
2. In the right pane, select **Control**.
3. Select the check box next to **Coh1**, and click **Start**.
4. Click **Yes** to confirm.

   The Node Manager starts the server on the target machine. When the Node Manager finishes its start sequence, the server's state is indicated in the State column in the Server Status table.
5. Repeat the above steps to start the remaining seven Coherence Servers.

5.13.9 Verify the Example

To verify the example:

Note: Before you verify the example, ensure that you start your managed server, as described in Section 5.16, "Starting Managed Servers on ComputeNode1 and ComputeNode2". In addition, copy `coherence.jar` to the `counter.war` file's `WEB-INF/lib` directory.

---

1. Open a browser and access the application, as in the following example:

   `http://`*`Managed_Server_host`*`:port/counter/counter.jsp`

   Managed_Server_host and *port* represent the host and the port of any Managed Server in `Dept1_Cluster1`. For example, the host and the port of `WLS1`.The `counter.war` application was deployed to this cluster.

   The counter page displays and the counter is set to 1, as illustrated in Figure 5–11.

Figure 5–11 Counter Page with Counter Set to 1

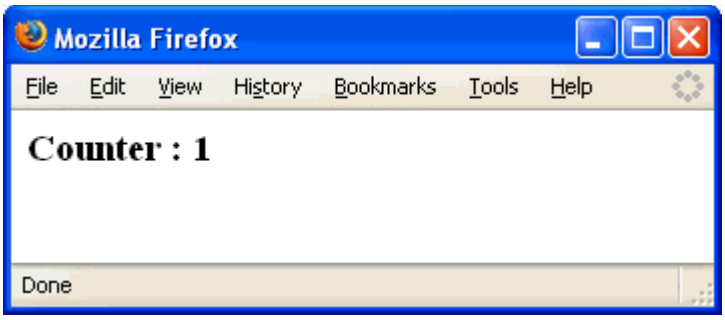

2. In a new browser (or new browser tab), access the `ServerB` counter instance using the following URL:

   `http://`*`Managed_Server_host`*`:port/counter/counter.jsp`

   Managed_Server_host and *port* represent the host and the port of a Managed Server in `Dept1_Cluster1`. For example, the host and the port of `WLS2`. The `WLS2` Managed Server should be listening on the same IP address as that of `WLS1`. However, the ports of `WLS1` and `WLS2` are different. The `counter.war` application was deployed to this cluster.

   The counter page is displayed, and the counter increments to 2 based on the session data, as shown in Figure 5–12.

Figure 5–12 Counter Page with Counter Set to 2

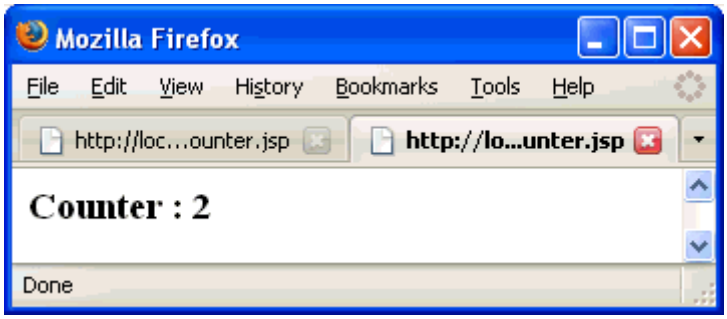

3. If you refresh the page, the counter increments to 3. Return to the original browser (or browser tab), refresh the instance, and the counter displays 4.

5.14 Specifying Node Manager Type for ComputeNode1 and ComputeNode2

You must specify the Node Manager type for `ComputeNode1` and `ComputeNode2` as follows:

1. Log in to the Oracle WebLogic Administration Console, using the following URL in a browser:

   `http://ADMINVHN1:7001/console`

2. From the **Domain Structure** menu, expand **Environment** and select **Machines**.

   The Summary of Machines page is displayed.

3. Select `ComputeNode1`. The Settings for ComputeNode1 page is displayed.

4. Click the Node Manager tab.
5. Ensure that `Plain` is selected as the **Type**. Click **Save**.
6. Repeat these steps for `ComputeNode2`.

5.15 Disabling Host Name Verification for Managed Servers

You will receive errors when managing the different WebLogic Servers if you have not configured the server certificates. To avoid these errors, disable host name verification while setting up and validating the topology.

To disable host name verification for Managed Servers, such as `WLS1`, complete the following steps:

1. Log in to Oracle WebLogic Server Administration Console.
2. Click **Lock & Edit**.
3. Expand the **Environment** node in the Domain Structure window.
4. Click **Servers**. The Summary of Servers page is displayed.
5. Select `WLS1` in the **Names** column of the table. The settings page for WLS1 is displayed.
6. Click the **SSL** tab.
7. Expand the **Advanced** section of the page.
8. Set host name verification to `None`.
9. Click **Save**.
10. Repeat steps 4 through 8 for the remaining Managed Servers in your WebLogic cluster.
11. Save and activate the changes.
12. Restart Node Manager:
    a. Stop Node Manager by stopping the process associated with it.

       If it is running in the foreground in a shell, simply use **Ctrl+c**.

       If it is running in the background in the shell, find the associate process and use the `kill` command to stop it.

    b. Start Node Manager.

5.16 Starting Managed Servers on ComputeNode1 and ComputeNode2

You must start the Managed Servers on `ComputeNode1` and `ComputeNode2` as follows:

1. Log in to the Oracle WebLogic Administration Console, using the following URL in a browser:

   ```
   http://ADMINVHN1:7001/console
   ```

2. From the **Domain Structure** menu, expand **Environment** and select **Servers**.

   The **Summary of Servers** page is displayed.
3. Select **Control**.
4. In the Servers table, click **WLS1** Managed Server instance.

5. Click **Start**.
6. On the **Server Life Cycle Assistant** page, click **Yes** to confirm.

   The Node Manager starts the server on the target machine. When the Node Manager finishes its start sequence, the server's state is indicated in the **State** column in the Servers Status table.
7. Repeat the above steps to start remaining Managed Servers on `ComputeNode1` and `ComputeNode2`.

5.17 Disabling Host Name Verification for the Administration Server

This step is required if you have not set up the appropriate certificates to authenticate the different nodes with the Administration Server. If you have not configured the server certificates, you will receive errors when managing the different WebLogic servers. To avoid these errors, disable host name verification while setting up and validating the topology.

Perform these steps to disable host name verification:

1. Log in to the Oracle WebLogic Administration Console, using the following URL in a browser:

   `http://ADMINVHN1:7001/console`
2. Click **Lock and Edit**.
3. From the **Domain Structure** menu, select **Environment**, and then **Servers**.

   The **Summary of Servers** page is displayed.
4. Select **AdminServer(admin)** in the Names column of the table. The settings page for the server is displayed.
5. Select the **SSL** tab.
6. Expand the **Advanced** section.
7. Set **Hostname Verification** to **None**.
8. Click **Save**.
9. Click **Activate Changes**.
10. The change will not take effect until the Administration Server is restarted (Node Manager must be up and running):
    a. In the Summary of Servers screen, select the **Control** tab.
    b. Select **AdminServer(admin)** in the table and then click **Shutdown**.
    c. Start the Administration Server again from the command line, as described in Section 5.6, "Starting the Administration Server on ComputeNode1."

5.18 Creating a JMS Persistence Store

Before you create the JMS persistence store, ensure that you have created a directory for the file store on your file system. Perform these steps to create a shared JMS persistence Store:

1. If you have not already done so, in the Change Center of the Administration Console, click **Lock & Edit**.

2. In the left pane of the Console, expand **Services** node and then select **Persistence Stores**.

   The **Summary of Persistence Stores** page is displayed.

3. Click **New**, and then **Create File Store**.

   Note: Once you create a file store, you cannot rename it. Instead, you must delete it and create another one that uses the new name.

4. On the Create a new File Store page, enter the following:

   - **Name**: Name for the File Store.
   - **Target**: WLS1.

     Note: In a clustered environment, a recommended best practice is to target a custom file store to the same migratable target as the migratable JMS service, so that a member server will not be a single point of failure. A file store can also be automatically migrated from an unhealthy server instance to a healthy server instance, with the help of the server health monitoring services. In addition, configure a reasonable maximum message count quota on each JMS server. Assume each message header uses approximately 1K.

   - **Directory**: Pathname to the directory on the file system where the file store is kept. This directory must exist on your system, so be sure to create it before completing this tab. Enter the location of the persistent storage on the Sun ZFS Storage 7320 appliance that is available to other servers in the cluster. Specifying this location enables pending JMS messages to be sent. The location should follow the following directory structure:

     `/u01/Dept_1/jmsjta/base_domain/Dept1_Cluster1/jms`

     Where `Dept1_Cluster1` is the name of the WebLogic cluster.

Note: This location maps to `/export/Dept_1/jmsjta` share on the Sun ZFS Storage 7320 appliance. Both `ComputeNode1` and `ComputeNode2` must be able to access this directory. Ensure that you mount this directory from both `ComputeNode1` and `ComputeNode2`, as described in Section 3.4.2, "Setting Up Enterprise Deployment Storage Configuration". This directory must also exist before you restart the server.

You must configure the location for all of the persistence stores as a directory that is visible from both `ComputeNode1` and `ComputeNode2`. Oracle recommends that you create and maintain separate shares for JMS logs under the `Dept_1` project, as shown in the Figure 3–5. See Section 3.4, "Shared Storage and Recommended Project and Share Structure" for more information.

You must change all of the persistent stores to use this shared base directory.

When a custom file store is targeted to a migratable target, the specified directory must be accessible from all candidate server members in the migratable target. For highest reliability, use a shared storage solution that is itself highly available—for example, a storage area network (SAN) or a dual-ported SCSI disk.

- Click **OK**.

5. Select the file store created in steps 2-3.
6. On the File Store, **Configuration** tab, update any additional Advanced parameters as required:
   - **Logical Name**: Optional name that can be used by subsystems that need a way to refer to different stores on different servers using the same name.
   - **Synchronous Write Policy**: Specifies how this file store writes data to disk. Choose **Direct-Write**, which is the default value.
7. Click **Save**.
8. Click **Activate Changes**.
9. Restart the servers to make the change in the persistent stores take effect.

Note: Each Managed Server, such as `WLS1`, gets a default file store.

If you are leveraging WebLogic JMS, Oracle recommends that you to follow additional configuration and tuning guidelines as described in: the following guides:

- "Tuning WebLogic JMS" in the *Oracle Fusion Middleware Performance and Tuning for Oracle WebLogic Server*
- "Best Practices for JMS Beginners and Advanced Users" in the *Oracle Fusion Middleware Configuring and Managing JMS for Oracle WebLogic Server*

5.19 Configuring a Default Persistence Store for Transaction Recovery

Each server has a transaction log that stores information about committed transactions that are coordinated by the server that may not have been completed. Oracle WebLogic Server uses this transaction log for recovery from system crashes or network failures. To leverage the migration capability of the Transaction Recovery Service for the servers within a cluster, store the transaction log in a location accessible to a server and its backup servers. In this example procedure, we describe how to configure a default persistence store for transaction recovery for the `WLS1` Managed Server to which the sample application is deployed. Follow the same procedure to configure a default persistence store for transaction recovery on the remaining Managed Servers, if required.

Note: This location should be on a shared file system the Sun ZFS Storage 7320 appliance.

Perform these steps to set the location for the default persistence store:

1. Log in to the Oracle WebLogic Server Administration Console.
2. In the left pane of the Console, expand **Environment** and select **Servers**.

   The **Summary of Servers** page is displayed.
3. Click the name of the server, such as `WLS1` (represented as a hyperlink) in the Name column of the table.

   The settings page for the selected server is displayed, and defaults to the Configuration tab.
4. Open the **Services** tab.
5. In the Default Store section of the page, enter the path to the folder where the default persistent stores will store its data files. The directory structure of the path is as follows:

   `/u01/Dept_1/jmsjta/base_domain/Dept1_Cluster1/tlogs`

   Where `Dept1_Cluster1` is the name of the WebLogic cluster.

   For more information, see Figure 3–5.
6. Click **Save**.

Note: To enable migration of the Transaction Recovery Service, specify a location on a persistent storage solution that is available to other servers in the cluster. Both `ComputeNode1` and `ComputeNode2` must be able to access this directory. Ensure that you mount this directory from both `ComputeNode1` and `ComputeNode2`, as described in Section 3.4.2, "Setting Up Enterprise Deployment Storage Configuration". This directory must also exist before you restart the server.

5.20 Manually Failing Over the Administration Server to ComputeNode2

In case the compute node (`ComputeNode1`) hosting the Administration Server fails, you can fail over the Administration Server to `ComputeNode2`. This section describes how to fail over the Administration Server from `ComputeNode1` to `ComputeNode2`.

Assumptions:

- The Administration Server is configured to listen on `10.0.0.17`. This address is the floating IP assigned to the Administration Server using the `BOND0` interface.
- The Administration Server is failed over from `ComputeNode1` to `ComputeNode2`, and the two nodes have these IPs (BOND0):
  - `ComputeNode1`: `192.168.10.1`
  - `ComputeNode2`: `192.168.10.2`
  - `10.0.0.17`: This is the floating IP where the Administration Server is running, assigned to `bond0:Y`, available in `ComputeNode1` and `ComputeNode2`.
- The domain directory where the Administration Server is running in `ComputeNode1` is on shared storage.

The following procedure shows how to fail over the Administration Server to a different node (`ComputeNode2`), but the Administration Server will still use the same WebLogic Server machine (which is a logical machine, not a physical machine).

1. Stop the Administration Server.
2. Migrate IP to the second node.
   a. Run the following command as root on `ComputeNode1` (where `bond0:Y` is the current interface used by `ADMINVHN1`):

      ```
      ComputeNode1> /sbin/ifconfig bond0:Y down
      ```

   b. Run the following command on `ComputeNode2`:

      ```
      ComputeNode2> /sbin/ifconfig <interface:index> <IP_Address> netmask
      <netmask>
      ```

      For example:

      ```
      /sbin/ifconfig bond0:1 10.0.0.17 netmask 255.255.255.0
      ```

      Note: Ensure that the netmask and interface to be used to match the available network configuration in `ComputeNode2`.

3. On Linux, update routing tables through `arping`, for example:

   ```
   ComputeNode2> /sbin/arping -b -A -c 3 -I bond0 10.0.0.17
   ```

4. Start the Administration Server on `ComputeNode2` using the `startWebLogic.sh` script, as described in Section 5.6, "Starting the Administration Server on ComputeNode1.".
5. Test that you can access the Administration Server on `ComputeNode2` as follows:
   a. Ensure that you can access the Oracle WebLogic Server Administration Console at `http://10.0.0.17:7001/console`.
   b. Check that you can access and verify the status of components in the Oracle Enterprise Manager at `http://ADMINVHN1:7001/em`.

Note: The Administration Server does not use Node Manager for failover. After a manual failover, the machine name that appears in the **Current Machine** field in the Administration Console for the server is `ComputeNode1`, and not the failover machine, `ComputeNode2`. Since Node Manager does not monitor the Administration Server, the machine name that appears in the **Current Machine** field, is not relevant and you can ignored it.

If you created a `boot.properties` file for the Administration Server on `ComputeNode1`, the username and password values in the file get encrypted. When the Administration Server is failed over to `ComputeNode2`, you must edit the username and password values in the `boot.properties` file on `ComputeNode2` manually.

5.21 Failing the Administration Server Back to ComputeNode1

This step checks that you can fail back the Administration Server, that is, stop it on `ComputeNode2` and run it on `ComputeNode1`. To do this, migrate Administration Server back to `ComputeNode1` as follows:

1. Run the following command on `ComputeNode2`.

   ```
   ComputeNode2> /sbin/ifconfig bond0:N down
   ```

2. Run the following command on `ComputeNode1`:

   ```
   ComputeNode1> /sbin/ifconfig bond0:Y 10.0.0.17 netmask 255.255.255.0
   ```

   Note: Ensure that the netmask and interface to be used match the available network configuration in `ComputeNode1`.

3. On Linux, update routing tables through `arping`. Run the following command from `ComputeNode1`.

   ```
   ComputeNode1> /sbin/arping -b -A -c 3 -I bond0 10.0.0.17
   ```

4. Start WLST and connect to Node Manager with **nmconnect** and start the Administration Server using **nmstart**.

   ```
   ComputeNode1> /u01/app/FMW_Product1/Oracle/Middleware/wlserver_10.3/common/bin
   ComputeNode1> ./wlst.sh
   ```

   Once you are in the `WLST` shell:

   ```
   wls:/offline>nmConnect('Admin_User','Admin_Password',
   'ComputeNode1','5556','domain_name','/u01/Dept_1/domains/el01cn01/base_
   domain','ssl')

   wls:/nm/base_domain> nmStart('AdminServer')
   ```

5. Test that you can access the Oracle WebLogic Server Administration Console at `http://10.0.0.17:7001/console`.

5.22 Backing Up Domain Configuration

Perform a backup to save your domain configuration (make sure that you stop the Administration Server and Managed Servers first). The configuration files all exist under the `/u01/Dept_1/domains/el01cn01/base_domain` directory.

```
tar -cvzf base_domainback.tar /u01/app/Dept_1/domain_home/base_domain
```

6 Configuring Oracle HTTP Server

This chapter describes how to configure Oracle HTTP Server 11*g* Release 1 (11.1.1.5.0) to support the Oracle Exalogic enterprise deployment.

To configure Oracle HTTP Server, you must complete the following steps:

- Section 6.1, "Important Notes Before You Begin"
- Section 6.2, "Prerequisites"
- Section 6.3, "Mandatory: Configuring Oracle HTTP Server for Administration Server and Managed Servers"
- Section 6.4, "Optional: Configuring Oracle HTTP Server for Load Balancing on the Private InfiniBand Network"
- Section 6.5, "Setting the Frontend URL for the Administration Console"
- Section 6.6, "Validating Access to ComputeNode1 Through Oracle HTTP Server"
- Section 6.7, "Validating Access to ComputeNode2 Through Oracle HTTP Server"

6.1 Important Notes Before You Begin

If you are an Oracle Solaris user, read Section 3.1, "Important Notes for Oracle Solaris Users" before you configure Oracle HTTP Server.

6.2 Prerequisites

The following are the prerequisites for configuring Oracle HTTP Server 11*g* Release 1 (11.1.1.5.0) for Oracle Exalogic:

- Ensure that you have your existing Oracle HTTP Server 11*g* set up, as described in Section 4.3, "Installing Oracle HTTP Server".
- It is assumed that Oracle HTTP Server 11*g* Release1 (11.1.1.5.0) is installed outside of the Oracle Exalogic environment.
- If you are using Oracle HTTP Server to load balance traffic on the `BOND0/IPMP0` network, ensure that Oracle HTTP Server is installed in the Exalogic environment, as described in Section 4.3, "Installing Oracle HTTP Server". In addition, you should have set up the required number of Oracle HTTP Server instances to run on Exalogic compute nodes, as described in Section 5.10, "Optional: Creating Oracle HTTP Server Instances in the Exalogic Environment".

6.3 Mandatory: Configuring Oracle HTTP Server for Administration Server and Managed Servers

You must complete this procedure to enable Oracle HTTP Server to route traffic to the Administration Server and to the Oracle WebLogic cluster . This Oracle HTTP Server installation is outside of Exalogic. `Dept1_Cluster1` is the cluster to which the application will be deployed, and you must set the `WebLogicCluster` parameter to the list of Managed Servers in the cluster as follows:

> **Note:** `WEBHOST1` and `WEBHOST2` are used as the example hosts for the Oracle HTTP Server instances (`OHS1` and `OHS2`) in these procedures.

1. On `WEBHOST1` and `WEBHOST2`, add the following lines to the *ORACLE_BASE/admin/<instance_name>/config/OHS/<component_name>/mod_wl_ohs.conf* file:

```
# Dept1_Cluster1
<Location /shopping-cart-webapp>
  SetHandler weblogic-handler
WebLogicCluster
10.1.0.1:7003,10.1.0.2:7003,10.1.0.3:7003,10.1.0.4:7003,10.1.0.9:7003,10.1.0.10
:7003,10.1.0.11:7003,10.1.0.12:7003
</Location>
```

> **Note:** In the `<Location /shopping-cart-webapp entry>`, ensure that you enter the relevant context paths to the web application that you will deploy.

2. Make sure that the *httpd.conf* file located in the same directory as the *mod_wl_ohs* file contains the following lines:

```
NameVirtualHost *:7777
<VirtualHost *:7777>
    ServerName https://exalogic.mycompany.com:443
    ServerAdmin you@your.address
    RewriteEngine On
    RewriteOptions inherit
</VirtualHost>

NameVirtualHost *:7777
<VirtualHost *:7777>
    ServerName http://exalogicinternal.mycompany.com:80
    ServerAdmin you@your.address
    RewriteEngine On
    RewriteOptions inherit
</VirtualHost>

NameVirtualHost *:7777
<VirtualHost *:7777>
    ServerName admin.mycompany.com:80
    ServerAdmin you@your.address
    RewriteEngine On
    RewriteOptions inherit
# Admin Server and EM
<Location /console>
```

```
    SetHandler weblogic-handler
    WebLogicHost 10.1.0.17
    WeblogicPort 7003
</Location>
<Location /consolehelp>
    SetHandler weblogic-handler
     WebLogicHost 10.1.0.17
    WeblogicPort 7003
 </Location>
<Location /em>
    SetHandler weblogic-handler
    WebLogicHost 10.1.0.17
    WeblogicPort 7003
</Location>
</VirtualHost>
```

Note: Values such as `exalogic.mycompany.com:443,7777`, `admin.mycompany:80`, and `you@youraddress` that are noted in this document serve as examples only. Enter values based on the actual environment.

The Administration Server uses a `BOND1` EoIB floating IP address, such as `10.1.0.17`. Access to the Administration Console is restricted to requests coming in from `admin.mycompany.com`.

3. Restart Oracle HTTP Server on both `WEBHOST1` and `WEBHOST2`.

```
WEBHOST1> ORACLE_BASE/admin/<instance_name>/bin/opmnctl restartproc
ias-component=ohs1

WEBHOST2> ORACLE_BASE/admin/<instance_name>/bin/opmnctl restartproc
ias-component=ohs2
```

The servers specified in the `WebLogicCluster` parameter are only important at startup time for the plug-in. The list needs to provide at least one running cluster member for the plug-in to discover other members of the cluster. Note that the listed cluster member must be running when Oracle HTTP Server is started. Oracle WebLogic Server and the plug-in work together to update the server list automatically with new, failed, and recovered cluster members.

The method of specifying the parameter, and the required format vary by plug-in. See the following examples in *Oracle Fusion Middleware Using Web Server 1.1 Plus-Ins with Oracle WebLogic Server:*

"Installing and Configuring the Apache HTTP Server Plug-In"

The plug-in does a simple round-robin between all available servers. The server list specified in this property is a starting point for the dynamic server list that the server and plug-in maintain. WebLogic Server and the plug-in work together to update the server list automatically with new, failed, and recovered cluster members.

The plug-in directs HTTP requests containing a cookie, URL-encoded session, or a session stored in the POST data to the server in the cluster that originally created the cookie.

Note: For more information about these parameters, see the "Parameters for Web Server Plug-Ins" topic in the *Oracle Fusion Middleware Using Web Server 1.1 Plus-Ins with Oracle WebLogic Server.*

6.4 Optional: Configuring Oracle HTTP Server for Load Balancing on the Private InfiniBand Network

You should complete this procedure if you are using Oracle HTTP Server to load balance traffic on the `BOND0/IPMP0` private network. This Oracle HTTP Server installation is inside the Exalogic environment. `Dept1_Cluster1` is the cluster to which the application will be deployed, and you must set the `WebLogicCluster` parameter to the list of Managed Servers in the cluster as follows:

Note: `ComputeNode1` and `ComputeNode2` are used as the example compute nodes that are hosting the Oracle HTTP Server instances (`OHS1` and `OHS2`) in these procedures.

1. On `ComputeNode1` and `ComputeNode2`, add the following lines to the `/u01/app/FMW_Product1/Oracle/Middleware/Oracle_WT1/instances/OHS1/config/OHS/ohs1/mod_wl_ohs.conf` file:

```
# Dept1_Cluster1
<Location /shopping-cart-webapp>
  SetHandler weblogic-handler
WebLogicCluster
10.0.0.1:7003,10.0.0.2:7003,10.0.0.3:7003,10.0.0.4:7003,10.0.0.9:7003,10.0.0.10
:7003,10.0.0.11:7003,10.0.0.12:7003
</Location>
```

Notes:

- In the `<Location /shopping-cart-webapp entry>`, ensure that you enter the relevant context paths to the web application that you will deploy.
- The above example includes entries for the cluster `Dept1_Cluster1`. These floating IP addresses use the `BOND0` interface via IP over InfiniBand (`IPoIB`).

2. Restart Oracle HTTP Server on both `ComputeNode1` and `ComputeNode2`, as in the following example:

```
# /u01/app/FMW_Product1/Oracle/Middleware/Oracle_WT1/opmn/bin/opmnctl restartproc ias-component=OHS1
```

   If you are using the second Oracle HTTP Server instance, such as `OHS2`, on the same compute node, restart both `OHS1` and `OHS2`.

The servers specified in the `WebLogicCluster` parameter are only important at startup time for the plug-in. The list needs to provide at least one running cluster member for the plug-in to discover other members of the cluster. Note that the listed cluster member must be running when Oracle HTTP Server is started. Oracle

WebLogic Server and the plug-in work together to update the server list automatically with new, failed, and recovered cluster members.

The method of specifying the parameter, and the required format vary by plug-in. See the following example in *Oracle Fusion Middleware Using Web Server 1.1 Plus-Ins with Oracle WebLogic Server:*

"Installing and Configuring the Apache HTTP Server Plug-In"

The plug-in does a simple round-robin between all available servers. The server list specified in this property is a starting point for the dynamic server list that the server and plug-in maintain. WebLogic Server and the plug-in work together to update the server list automatically with new, failed, and recovered cluster members.

The plug-in directs HTTP requests containing a cookie, URL-encoded session, or a session stored in the POST data to the server in the cluster that originally created the cookie.

Note: For more information about these parameters, see the "Parameters for Web Server Plug-Ins" topic in the *Oracle Fusion Middleware Using Web Server 1.1 Plus-Ins with Oracle WebLogic Server*.

6.5 Setting the Frontend URL for the Administration Console

The Oracle WebLogic Server Administration Console application tracks changes made to ports, channels and security using the console. When changes made through the console are activated, the console validates its current listen address, port and protocol. If the listen address, port and protocol are still valid, the console redirects the HTTP request replacing the host and port information with the Administration Server's listen address and port. When the Administration Console is accessed using a load balancing router (LBR), it is required to change the Administration Server's frontend URL, so that the user's web browser is redirected to the appropriate LBR address. To do this, complete these steps:

1. Log in to the Oracle WebLogic Server Administration Console.
2. Click **Lock and Edit**.
3. From the **Domain Structure** menu, select **Environment**, and then **Servers**.

   The **Summary of Servers** page is displayed.
4. Select **Admin Server** in the Names column of the table.

   The settings page for AdminServer(admin) is displayed.
5. Click the **Protocols** tab.
6. Click the **HTTP** tab.
7. Set the **Front End Host** field to `admin.mycompany.com` (your LBR address).
8. Save and activate the changes.

To eliminate redirections, it is recommended that you disable the Administration Console's "Follow changes" feature. To do this, log in to the Administration Console, click **Preferences** and then **Shared Preferences**. Clear the '**Follow Configuration Changes**' check box, and click **Save**.

6.6 Validating Access to ComputeNode1 Through Oracle HTTP Server

Verify that the server status is reported as "Running" in the Administration Console. If the server is shown as "Starting" or "Resuming," wait for the server status to change to "Started." If another status is reported (such as "Admin" or "Failed"), check the server output log files for errors.

Validate Administration Console and Enterprise Manager through both Oracle HTTP Server instances using the following URLs:

- `http://WEBHOST1:7777/console`
- `http://WEBHOST2:7777/console`
- `http://admin.mycompany.com/console`

For information on configuring system access through the load balancer, see Section 3.3.5, "Load Balancers."

6.7 Validating Access to ComputeNode2 Through Oracle HTTP Server

Perform the same steps as in Section 6.6, "Validating Access to ComputeNode1 Through Oracle HTTP Server". This is to check that you can access the Administration Server when it is running on `ComputeNode2`.

7

Enabling Exalogic-Specific Enhancements in Oracle WebLogic Server 11*g* Release 1 (10.3.4)

This section discusses the following topics:

7.1 Important Notes Before You Begin

If you are an Oracle Solaris user, read Section 3.1, "Important Notes for Oracle Solaris Users" before you complete the procedures described in this chapter.

7.2 Overview of Exalogic-Specific Enhancements

Oracle Exalogic includes performance optimizations for Oracle WebLogic Server to improve input/output, thread management, and request handling efficiency. A WebLogic Server domain can be configured to enable domain-wide input/output optimizations. These optimizations include multi-core architectural enhancements that improve thread management, request processing, and reduce lock contention.

Additional optimizations include reduced buffer copies, which result in more efficient input/output. Finally, session replication performance and CPU utilization is improved through lazy deserialization, which avoids performing extra work on every session update that is only necessary when a server fails.

WebLogic Server clusters can be configured with cluster-wide optimizations that further improve server-to-server communication. The first optimization enables multiple replication channels, which improve network throughput among WebLogic Server cluster nodes. The second cluster optimization enables InfiniBand support for Sockets Direct Protocol, which reduces CPU utilization as network traffic bypasses the TCP stack.

7.3 Prerequisites

The following are the prerequisites for configuring Oracle Fusion Middleware 11*g* Release 1 (11.1.1) products for Oracle Exalogic:

- Preconfiguring the environment, including database, storage, and network, as described in Chapter 3, "Network, Storage, and Database Preconfiguration".
- Your Oracle Exalogic Domain is configured, as described in Chapter 5, "Configuring Oracle Fusion Middleware".

7.4 Enabling Domain-Level Enhancements

To enable domain-level enhancements, complete the following steps:

1. Log in to the Oracle WebLogic Server Administration Console.
2. Select *Domainname* in the left navigation pane. The Settings for *Domainname* screen is displayed. Click the **General** tab.
3. In your domain home page, select **Enable Exalogic Optimizations**, and click **Save**.
4. Activate changes.
5. Stop and start your domain.

The **Enable Exalogic Optimizations** setting collectively enables all of the individual features described in Table 7–1. The Startup Option column indicates how to independently enable and disable each fe ature.

Table 7–1 Features Enabled by the Domain-Level Flag

Feature	Description	Startup Option	MBean
Scattered Reads	Increased efficiency during I/O in environments with high network throughput	`-Dweblogic.ScatteredReadsEnabled=true/false`	`KernelMBean.setScatteredReadsEnabled`
Gathered Writes	Increased efficiency during I/O in environments with high network throughput	`-Dweblogic.GatheredWritesEnabled=true/false`	`KernelMBean.setGatheredWritesEnabled`
Lazy Deserialization	Increased efficiency with session replication	`-Dweblogic.replication.enableLazyDeserialization=true/false`	`ClusterMBean.setSessionLazyDeserializationEnabled`
Self Tuning Thread Pool Optimization	Increased efficiency of the self tuning thread pool by aligning it with the Exalogic's processor architecture threading capabilities	Not applicable	`KernelMBean.addWorkManagerThreadsByCpuCount`

Note: After enabling the optimizations, you may see the following message:

```
java.io.IOException: Broken pipe
```

You may see the same message when storage failover occurs. In either case, you can ignore the error message.

7.5 Enabling Cluster-Level Session Replication Enhancements

You can enable session replication enhancements for Managed Servers in a WebLogic cluster to which you will deploy a web application at a later time.

Note: If you are using Coherence*web, these session replication enhancements do not apply.

If you use the `dizzyworld.ear` application as described in Chapter 8, "Deploying a Sample Web Application to an Oracle WebLogic Cluster" then you should skip these steps.

To enable session replication enhancements for `Dept1_Cluster1`, complete the following steps:

1. Ensure that Managed Servers in the `Dept1_Cluster1` cluster are up and running, as described in Section 5.16, "Starting Managed Servers on ComputeNode1 and ComputeNode2".
2. To set replication ports for a Managed Server, such as `WLS1`, complete the following steps:
   a. Under **Domain Structure**, click **Environment** and **Servers**. The Summary of Servers page is displayed.
   b. Click `WLS1` on the list of servers. The Settings for WLS1 is displayed.
   c. Click the **Cluster** tab.
   d. In the **Replication Ports** field, enter a range of ports for configuring multiple replication channels. For example, replication channels for Managed Servers in `Dept_1_Cluster1` can listen on ports starting from `7005` to `7015`. To specify this range of ports, enter `7005-7015`.
3. Create a custom network channel for each Managed Server in the cluster (for example, `WLS1`) as follows:
   a. Log in to the Oracle WebLogic Server Administration Console.
   b. If you have not already done so, click **Lock & Edit** in the Change Center.
   c. In the left pane of the Console, expand **Environment** and select **Servers**.

      The **Summary of Servers page** is displayed.
   d. In the Servers table, click **WLS1** Managed Server instance.
   e. Select **Protocols**, and then **Channels**.
   f. Click **New**.
   g. Enter **ReplicationChannel** as the name of the new network channel and select **t3** as the protocol, then click **Next**.
   h. Enter the following information:

      Listen address: **10.0.0.1**

      Note: This is the floating IP assigned to `WLS1`.

      Listen port: **7005**

i. Click **Next**, and in the Network Channel Properties page, select **Enabled** and **Outbound Enabled**.

j. Click **Finish**.

k. Under the Network Channels table, select **ReplicationChannel**, the network channel you created for the `WLS1` Managed Server.

l. Expand **Advanced**, and select **Enable SDP Protocol**.

m. Click **Save**.

n. To activate these changes, in the Change Center of the Administration Console, click **Activate Changes**.

You must repeat the above steps to create a network channel each for the remaining Managed Servers in the `Dept1_Cluster1` cluster. Enter the required properties, as described in Table 7–2.

Table 7–2 Network Channels Properties

Managed Servers in Dept1_ Cluster1	Name	Protocol	Listen Address	Listen Port	Additional Channel Ports
WLS2	`ReplicationChannel`	t3	10.0.0.2	7005	7006 to 7014
WLS3	`ReplicationChannel`	t3	10.0.0.3	7005	7006 to 7014
WLS4	`ReplicationChannel`	t3	10.0.0.4	7005	7006 to 7014
WLS5	`ReplicationChannel`	t3	10.0.0.5	7005	7006 to 7014
WLS6	`ReplicationChannel`	t3	10.0.0.6	7005	7006 to 7014
WLS7	`ReplicationChannel`	t3	10.0.0.7	7005	7006 to 7014
WLS8	`ReplicationChannel`	t3	10.0.0.8	7005	7006 to 7014

4. After creating the network channel for each of the Managed Servers in your cluster, click **Environment** > **Clusters**. The Summary of Clusters page is displayed.

5. Click `Dept1_Cluster1` (this is the example cluster to which you will deploy a web application at a later time). The Settings for Dept1_Cluster1 page is displayed.

6. Click the **Replication** tab.

7. In the **Replication Channel** field, ensure that `ReplicationChannel` is set as the name of the channel to be used for replication traffic.

8. In the **Advanced** section, select the **Enable One Way RMI for Replication** option.

9. Click **Save**.

10. Activate changes, and restart the Managed Servers.

11. Manually add the system property `-Djava.net.preferIPv4Stack=true` to the **startWebLogic.sh** script, which is located in the `bin` directory of `base_domain`, using a text editor as follows:

    a. Locate the following line in the startWebLogic.sh script:

    ```
    . ${DOMAIN_HOME}/bin/setDomainEnv.sh $*
    ```

    b. Add the following property immediately after the above entry:

```
JAVA_OPTIONS="${JAVA_OPTIONS}
-Djava.net.preferIPv4Stack=true"
```

c. Save the file and close.

12. Restart all Managed Servers as follows:

a. In the administration console, click **Environment** > **Servers**. The Summary of Servers page is displayed.

b. Select a Managed Server, such as `WLS1`, by clicking `WLS1`. The Settings for WLS1 page is displayed.

c. Click the **Control** tab. Select `WLS1` in the Server Status table. Click **Start**.

d. Repeat these steps for each of the Managed Servers in the WebLogic cluster.

Note: To verify that multiple listening ports were opened, you can either run the `netstat -na` command on the command line or check the Managed Server logs.

7.6 Configuring Grid Link Data Source for Dept1_Cluster1

You must create a Grid Link Data Source for JDBC connectivity between Oracle WebLogic Server and a service targeted to an Oracle RAC cluster. It uses the Oracle Notification Service (ONS) to adaptively respond to state changes in an Oracle RAC instance. This section includes the following:

- What is a Grid Link Data Source
- Creating a GridLink Data Source on Dept1_Cluster1

7.6.1 What is a Grid Link Data Source

A Grid Link data source includes the features of generic data sources plus the following support for Oracle RAC:

- Section 7.6.1.1, "Fast Connection Failover"
- Section 7.6.1.2, "Runtime Connection Load Balancing"
- Section 7.6.1.3, "XA Affinity"
- Section 7.6.1.4, "SCAN Addresses"
- Section 7.6.1.5, "Secure Communication using Oracle Wallet."

7.6.1.1 Fast Connection Failover

A Grid Link data source uses Fast Connection Failover to:

- Provide rapid failure detection.
- Abort and remove invalid connections from the connection pool.
- Perform graceful shutdown for planned and unplanned Oracle RAC node outages. The data source allows in-progress transactions to complete before closing connections. New requests are load balanced to an active Oracle RAC node.
- Adapt to changes in topology, such as adding a new node.
- Distribute runtime work requests to all active Oracle RAC instances.

See Fast Connection Failover in the *Oracle Database JDBC Developer's Guide and Reference*.

7.6.1.2 Runtime Connection Load Balancing

Runtime Connection Load Balancing allows WebLogic Server to:

- Adjust the distribution of work based on back end node capacities such as CPU, availability, and response time.
- React to changes in Oracle RAC topology.
- Manage pooled connections for high performance and scalability.

If FAN is not enabled, Grid link data sources use a round-robin load balancing algorithm to allocate connections to Oracle RAC nodes.

7.6.1.3 XA Affinity

XA Affinity for global transactions ensures all the data base operations for a global transaction performed on an Oracle RAC cluster are directed to the same Oracle RAC instance. The first connection request for an XA transaction is load balanced using RCLB and is assigned an Affinity context. All subsequent connection requests are routed to the same Oracle RAC instance using the Affinity context of the first connection.

7.6.1.4 SCAN Addresses

Oracle Single Client Access Name (SCAN) addresses can be used to specify the host and port for both the TNS listener and the ONS listener in the WebLogic console. A Grid Link data source containing SCAN addresses does not need to change if you add or remove Oracle RAC nodes. Contact your network administrator for appropriately configured SCAN urls for your environment. For more information, see http://www.oracle.com/technetwork/database/clustering/overview/scan-129069.pdf.

7.6.1.5 Secure Communication using Oracle Wallet

Allows you to configure secure communication with the ONS listener using Oracle Wallet.

7.6.2 Creating a GridLink Data Source on Dept1_Cluster1

You create a Grid Link data source for each of the Oracle database instances, both for these data sources and the global leasing data source. When you create a data source:

- Make sure that this is a non-xa data source (This applies only for the global leasing data source).
- Target these data sources to the `Dept1_Cluster1` cluster.
- Make sure that the datasources connection pool initial capacity is set to **0**. To do this, in the Oracle WebLogic Server Administration Console, select **Services**, **JDBC**, and then **Datasources**. In the Datasources screen, click the **Datasource Name**, then click the **Connection Pool** tab, and enter **0** in the **Initial capacity** field.
- Ensure that an ONS daemon is running on your database servers at all times. You can start the ONS daemon on a database server by running the `onsctl` command:

  ```
  start
  ```

To create a Grid Link Data source, complete the following steps:

1. Log in to the Oracle WebLogic Server Administration Console.
2. If you have not already done so, in the Change Center of the Administration Console, click **Lock & Edit**.
3. In the **Domain Structure** tree, expand **Services**, then select **Data Sources**.
4. On the **Summary of Data Sources** page, click **New** and select **GridLink Data Source**.

   The **Create a New JDBC GridLink Data Source** page is displayed.
5. Enter the following information:
   - Enter a logical name for the datasource in the Name field. For example, **gridlink**.
   - Enter a name for JNDI. For example, **jdbc/gridlink**.
   - Click **Next.**
6. In the **Transaction Options** page, de-select **Supports Global Transactions**, and click **Next**.
7. Select **Enter individual listener information** and click **Next**.
8. Enter the following connection properties:
   - **Service Name**: Enter the name of the Oracle RAC service in the **Service Name** field. For example, enter **myService** in **Service Name**.

     Note: The Oracle RAC Service name is defined on the database, and it is not a fixed name.

   - **Host Name** - Enter the DNS name or IP address of the server that hosts the database. For an Oracle GridLink service-instance connection, this must be the same for each data source in a given multi data source.
   - **Port** - Enter the port on which the database server listens for connections requests.
   - **Database User Name**: Enter the database user name. For example, **myDataBase**.
   - Password: Enter the password. For example, **myPassword1**.
   - Confirm Password and click **Next**.

     Tip: For more information, see the *Oracle Fusion Middleware Oracle WebLogic Server Administration Console Online Help*.

   The console automatically generates the complete JDBC URL. For example:

   ```
   jdbc:oracle:thin:@(DESCRIPTION=(ADDRESS_
   LIST=(ADDRESS=(PROTOCOL=TCP)(HOST=left)(PORT=1234))(ADDRESS=(PROTOCOL=TCP)(HOST
   =right)(PORT=1234))(ADDRESS=(PROTOCOL=TCP)(HOST=center)(PORT=1234)))(CONNECT_
   DATA=(SERVICE_NAME=myService)))
   ```
9. On the **Test GridLink Database Connection** page, review the connection parameters and click **Test All Listeners**.

   Oracle WebLogic attempts to create a connection from the Administration Server to the database. Results from the connection test are displayed at the top of the

page. If the test is unsuccessful, you should correct any configuration errors and retry the test.

Click **Next**.

10. In the **ONS Client Configuration** page, do the following:
    - Select **Fan Enabled** to subscribe to and process Oracle FAN events.
    - In ONS host and port, enter a comma-separate list of ONS daemon listen addresses and ports for receiving ONS-based FAN events. You can use Single Client Access Name (SCAN) addresses to access FAN notifications.
    - Click **Next**.
11. On the **Test ONS client configuration** page, review the connection parameters and click **Test All ONS Nodes**.

    Click **Next**.
12. In the **Select Targets** page, select **Dept1_Cluster1** as the target and **All Servers in the cluster**
13. Click **Finish**.
14. Click **Activate Changes**.

7.7 Configuring SDP-Enabled JDBC Drivers for Dept1_Cluster1

You must configure SDP-enabled JDBC drivers for the `Dept1_Cluster1` cluster.

This section discusses the following topics:

7.7.1 Prerequisite

Before enabling SDP support for JDBC, you must configure the database to support InfiniBand, as described in the "Configuring SDP Protocol Support for Infiniband Network Communication to the Database Server" topic in the *Oracle Database Net Services Administrator's Guide*. Ensure that you set the protocol to SDP.

7.7.2 Enabling SDP Support for JDBC

Complete the following steps:

1. Ensure that you have created the Grid Link Data Sources for the JDBC connectivity on `ComputeNode1` and `ComputeNode2`, as described in Section 7.6, "Configuring Grid Link Data Source for Dept1_Cluster1".

   The console automatically generates the complete JDBC URL, as shown in the following example:

   ```
   jdbc:oracle:thin:@(DESCRIPTION=(ADDRESS=(PROTOCOL=tcp)(HOST=192.x.x.x)(PORT=152
   2))(CONNECT_DATA=(SERVICE_NAME=myservice)))
   ```

2. In the JDBC URL, replace `TCP` protocol with `SDP` protocol. For example:

   ```
   jdbc:oracle:thin:@(DESCRIPTION=(ADDRESS=(PROTOCOL=sdp)(HOST=192.x.x.x)(PORT=152
   ```

```
2))(CONNECT_DATA=(SERVICE_NAME=myservice)))
```

3. Manually add the system property `-Djava.net.preferIPv4Stack=true` to the **startWebLogic.sh** script, which is located in the `bin` directory of `base_domain`, using a text editor as follows:
   a. Locate the following line in the `startWebLogic.sh` script:
      ```
      . ${DOMAIN_HOME}/bin/setDomainEnv.sh $*
      ```
   b. Add the following property immediately after the above entry:
      ```
      JAVA_OPTIONS="${JAVA_OPTIONS} -Djava.net.preferIPv4Stack=true
      -Doracle.net.SDP=true"
      ```
   c. Save the file and close.
4. Restart the Managed Server as follows:
   a. In the administration console, click **Environment** > **Servers**. The Summary of Servers page is displayed.
   b. Select a Managed Server, such as `WLS1`, by clicking `WLS1`. The Settings for WLS1 page is displayed.
   c. Click the **Control** tab. Select `WLS1` in the Server Status table. Click **Start**.

7.7.3 Monitoring SDP Sockets Using sdpnetstat on Oracle Linux

You can monitor SDP sockets by running the `sdpnetstat` command on the Oracle Linux operating system installed on Exalogic compute nodes. If you have connected your Exalogic machine to the Oracle Exadata Database Machine, you should run the sdpnetstat command on both Exalogic compute nodes and Oracle Exadata database servers to monitor SDP traffic between your Exalogic machine and the Oracle Exadata Database Machine.

Log in to the operating system as a `root`, and run the following command on the command line:

```
# sdpnetstat
```

This command displays the status of SDP sockets (`ESTABLISHED` or not), as in the following sample output:

```
Proto  Recv-Q  Send-Q  Local Address         Foreign Address      State
tcp    0   0  l01cn04-priv.local:844 192.168.10.97:nfs   ESTABLISHED
sdp    0   0  defg04:5901       abc-pc.cd.mno:13982  ESTABLISHED
tcp    0   0  defg04:58192       defg.hg.test.:5901  ESTABLISHED
sdp    0   104  defg04:ssh       dhcp-loc-bldg:4706  ESTABLISHED
tcp    0   52  defg04:ssh       dhcp-loc-bldg-1:gsakmp  ESTABLISHED
```

7.7.4 Monitoring SDP Sockets Using netstat on Oracle Solaris

If you are using the Oracle Solaris operating system on your Exalogic compute nodes, you can run the following command after logging in as `root`:

```
# netstat -f sdp
```

The output of this command displays the address state for an SDP socket.

7.8 Configuring SDP InfiniBand Listener for Exalogic Connections

This section is intended for Exalogic users who have connected an Exalogic machine to an Oracle Exadata Database Machine.

In this scenario, database connections must use the InfiniBand network. Engineers who are deploying the Exalogic machine and the Oracle Exadata Database Machine for a multi-rack setup must coordinate the allocation of the IP addresses to ensure that both systems use the same subnet for the Infiniband network. In addition, IP addresses used on each system must be unique.

This section describes how to create an SDP listener on the InfiniBand network. The tasks described in this section are completed on the database nodes in the Oracle Exadata Database Machine running Oracle Linux.

Configuring the SDP listener involves the following tasks:

1. Enabling SDP on Database Nodes
2. Creating an SDP Listener on the InfiniBand Network

7.8.1 Enabling SDP on Database Nodes

To enable SDP on database nodes in the Oracle Exadata Database Machine running Oracle Linux, complete the following steps:

1. Open `/etc/infiniband/openib.conf` file in a text editor, and add the the following:

   ```
   set: SDP_LOAD=yes
   ```

2. Save the file and close.
3. Open the `/etc/ofed/libsdp.conf` file in a text editor, and edit the file as follows:

   To use both SDP and TCP, add the `use both` rule as follows:

   ```
   use both server * :
   use both client * :
   ```

   To exclude SDP (that is, to use only TCP), add the `use tcp` rule as follows:

   ```
   use tcp server * *:*
   use tcp client * *:*
   ```

4. Save the file and close.
5. Open `/etc/modprobe.conf` file in a text editor, and add the following setting:

   ```
   options ib_sdp sdp_zcopy_thresh=0 recv_poll=0
   ```

6. Save the file and close.
7. Reboot all database nodes for the changes to take effect.

7.8.2 Creating an SDP Listener on the InfiniBand Network

Oracle RAC 11*g* Release 2 supports client connections across multiple networks, and it provides load balancing and failover of client connections within the network they are connecting. To add a listener for the Exalogic connections coming in on the Infiniband network, first add a network resource for the infiniband network with Virtual IP addresses.

Note: This example lists two nodes (Oracle Exadata Database Machine quarter rack). If you have an Oracle Exadata Database Machine half or full rack, you must repeat node-specific lines for each node in the cluster.

1. Edit `/etc/hosts` on each node in the cluster to add the virtual IP addresses you will use for the InfiniBand network. Make sure that these IP addresses are not used. The following is an example:

```
# Added for Listener over IB
192.168.10.21 dm01db01-ibvip.mycompany.com dm01db01-ibvip
192.168.10.22 dm01db02-ibvip.mycompany.com dm01db02-ibvip
```

2. On one of the database nodes, as the `root` user, to create a network resource for the InfiniBand network, as in the following example:

```
# /u01/app/grid/product/11.2.0.2/bin/srvctl add network -k 2 -S
192.168.10.0/255.255.255.0/bondib0
```

3. Validate that the network was added correctly, by running one of the following commands:

```
# /u01/app/grid/product/11.2.0.2/bin/crsctl stat res -t | grep net
ora.net1.network
ora.net2.network -- Output indicating new Network resource
```

or

```
# /u01/app/grid/product/11.2.0.2/bin/srvctl config network -k 2
Network exists: 2/192.168.10.0/255.255.255.0/bondib0, type static -- Output
indicating Network resource on the 192.168.10.0 subnet
```

4. Add the Virtual IP addresses on the network created in Step 2, for each node in the cluster.

```
srvctl add vip -n dm01db01 -A dm01db01-ibvip/255.255.255.0/bondib0 -k 2
srvctl add vip -n dm01db02 -A dm01db02-ibvip/255.255.255.0/bondib0 -k 2
```

5. As the "`oracle`" user (who owns the Grid Infrastructure Home), add a listener which will listen on the VIP addresses created in Step 3.

```
srvctl add listener -l LISTENER_IB -k 2 -p TCP:1522,/SDP:1522
```

6. For each database that will accept connections from the middle tier, modify the `listener_networks init` parameter to allow load balancing and failover across multiple networks (Ethernet and InfiniBand). You can either enter the full tnsnames syntax in the initialization parameter or create entries in `tnsnames.ora` in `$ORACLE_HOME/network/admin` directory. The `TNSNAMES.ORA` entries must exist in the `GRID_HOME`. The following example first updates `tnsnames.ora`. Complete this step on each node in the cluster with the correct IP addresses for that node. `LISTENER_IBREMOTE` should list all other nodes that are in the cluster. `DBM_IB` should list all nodes in the cluster.

Note: The TNSNAMES entry is only read by the database instance on startup, if you modify the entry that is referred to by any init.ora parameter (LISTENER_NETWORKS), you must restart the instance or issue an ALTER SYSTEM SET LISTENER_NETWORKS command for the modifications to take affect by the instance.

```
DBM =
(DESCRIPTION =
(ADDRESS = (PROTOCOL = TCP)(HOST = dm01-scan)(PORT = 1521))
(CONNECT_DATA =
(SERVER = DEDICATED)
(SERVICE_NAME = dbm)
))

DBM_IB =
(DESCRIPTION =
(LOAD_BALANCE=on)
(ADDRESS = (PROTOCOL = TCP)(HOST = dm01db01-ibvip)(PORT = 1522))
(ADDRESS = (PROTOCOL = TCP)(HOST = dm01db02-ibvip)(PORT = 1522))
(CONNECT_DATA =
(SERVER = DEDICATED)
(SERVICE_NAME = dbm)
))

LISTENER_IBREMOTE =
(DESCRIPTION =
(ADDRESS_LIST =
(ADDRESS = (PROTOCOL = TCP)(HOST = dm01db02-ibvip.mycompany.com)(PORT = 1522))
))

LISTENER_IBLOCAL =
(DESCRIPTION =
(ADDRESS_LIST =
(ADDRESS = (PROTOCOL = TCP)(HOST = dm01db01-ibvip.mycompany.com)(PORT = 1522))
(ADDRESS = (PROTOCOL = SDP)(HOST = dm01db01-ibvip.mycompany.com)(PORT = 1522))
))

LISTENER_IPLOCAL =
(DESCRIPTION =
(ADDRESS_LIST =
(ADDRESS = (PROTOCOL = TCP)(HOST = dm0101-vip.mycompany.com)(PORT = 1521))
))

LISTENER_IPREMOTE =
(DESCRIPTION =
(ADDRESS_LIST =
(ADDRESS = (PROTOCOL = TCP)(HOST = dm01-scan.mycompany.com)(PORT = 1521))
))
```

7. Modify the listener_networks init parameter. Connect to the database instance as sysdba.

```
SQLPLUS> alter system set listener_networks='((NAME=network2) (LOCAL_
LISTENER=LISTENER_IBLOCAL)(REMOTE_LISTENER=LISTENER_IBREMOTE))',
'((NAME=network1)(LOCAL_LISTENER=LISTENER_IPLOCAL)(REMOTE_LISTENER=LISTENER_
IPREMOTE))' scope=both;
```

8. Stop and start LISTENER_IB for the modification in Step 7.

```
srvctl stop listener -l LISTENER_IB

srvctl start listener -l LISTENER_IB
```

8

Deploying a Sample Web Application to an Oracle WebLogic Cluster

In this example, you can deploy a sample application named `dizzyworld.ear` to the Managed Servers in the `Dept1_Cluster1` cluster.

> **Note:** This chapter discusses how to deploy the sample application to a single cluster (`Dept1_Cluster1`). You can choose to deploy the whole application in a single cluster or its different tiers in multiple clusters. In addition, you can deploy a separate application to a separate cluster within the same domain, based on your specific deployment and management requirements.

This chapter discusses the following topics:

8.1 Downloading the dizzyworld.ear Application

To download and extract the sample application, do the following:

1. Download the `dizzyworld.ear` file from the following URL:

   `https://www.samplecode.oracle.com/tracker/tracking/linkid/prpl1004/remcurreport/true/template/ViewIssue.vm/id/S587/nbrresults/111`

   > **Note:** If this file is downloaded as `dizzyworld.zip` in your web browser, manually rename its file name to `dizzyworld.ear`.

2. Copy the **dizzyworld.ear** file to a local directory in `/u01/common/general`.

   > **Note:** If you are deploying this sample application for testing purposes, you do not need to perform server migration.

8.2 Configuring WAR-Scoped Coherence Clusters

With this configuration, or if you want only one application to use Coherence caches, each deployed Web application becomes its own Coherence cluster. Caches will be visible to the individual modules only. For example, this could be a recommended deployment for a stand-alone WAR deployment or stand-alone EJB deployment.

If you are deploying multiple WAR files, note that this configuration produces the largest number of Coherence nodes in the cluster—one for each deployed WAR file that uses `coherence.jar`. It also results in the largest resource utilization of the three configurations—one copy of the Coherence classes are loaded for each deployed WAR. On the other hand, since each deployed Web application is its own cluster, Web applications are completely isolated from other potentially misbehaving Web applications.

Note: A Web module within an EAR can have a module-scoped Coherence node but an EJB module within an EAR can only have an application-scoped Coherence node.

To Use Coherence Caches with WAR-Scoped Clusters

1. Use the WebLogic Server Administration Console to deploy `coherence.jar` and `active-cache.jar` as shared libraries to all of the target servers where the application will be deployed. See "Install a Java EE Library" in the *Oracle Fusion Middleware Oracle WebLogic Server Administration Console Help*.

   As an alternative to the Administration Console, you can also deploy the JAR files on the command line. The following are sample deployment commands:

   ```
   java weblogic.Deployer -username <> -password <> -adminurl <> -deploy
   coherence.jar -name coherence -library -targets <>

   java weblogic.Deployer -username <> -password <> -adminurl <> -deploy
   active-cache.jar -name active-cache -library -targets <>
   ```

2. Import `coherence.jar` and `active-cache.jar` as optional packages in the `manifest.mf` file of each module that will be using Coherence.

   As an alternative to using the manifest file, copy `coherence.jar` and `active-cache.jar` to each WAR file's `WEB-INF/lib` directory.

   Example 8–1 illustrates the contents of a sample `manifest.mf` file.

Example 8–1 Referencing coherence and active-cache JAR Files in the manifest.mf File

```
Manifest-Version: 1.0
Extension-List: coherence active-cache
coherence-Extension-Name: coherence
active-cache-Extension-Name: active-cache
```

3. (Optional) If you want to configure Coherence cluster properties, create a `CoherenceClusterSystemResourceMBean` and reference it as a `coherence-cluster-ref` element in `weblogic.xml` or `weblogic-ejb-jar.xml` file.

   Example 8–2 illustrates a sample configuration for WAR-scoped clusters in the `weblogic.xml` file. The `myCoherenceCluster` MBean is of type `CoherenceClusterSystemResourceMBean`.

Example 8–2 coherence-cluster-ref Element for WAR-Scoped Clusters

```
<weblogic-web-app>
...
  <coherence-cluster-ref>
    <coherence-cluster-name>
     CoherenceCluster1
    </coherence-cluster-name>
  </coherence-cluster-ref>
...
</weblogic-web-app>
```

8.3 Installing and Deploying an Enterprise Application

For this enterprise deployment topology, you must install and deploy the `dizzyworld.ear` web application to Managed Servers in `Dept1_Cluster1`. To do so, compete the following steps:

1. Extract the **dizzyworld.ear** file to a local directory in `/u01/common/general`.
2. Open the `weblogic-application.xml` (located at `/u01/app/FMW_Product1/Oracle/Middleware/wlserver_10.3/server/lib/consoleapp/META-INF`) in a text editor and enter your Oracle Coherence cluster name, as shown in the following example:

   ```
   <wls:coherence-cluster-ref>
   <wls:coherence-cluster-name>CoherenceCluster1</wls:coherence-cluster-name>
   ```

   Save the file after making changes.
3. Rebuild the `dizzyworld.ear` application.
4. Log in to the WebLogic Administration Console.
5. If you have not already done so, in the Change Center of the Administration Console, click **Lock & Edit**.
6. In the left pane of the console, select **Deployments**.

   The **Summary of Deployments** page is displayed.
7. In the right pane, click **Install**.

   The **Install Application Assistant** page is displayed.
8. Specify the path, where your `dizzyworld.ear` file is located, and click **Next**.
9. Select **Install this deployment as an application** and click **Next**.
10. Select **Dept1_Cluster1**, and then **All servers in the cluster** (Figure 8–1).

Figure 8–1 Select Deployment Targets

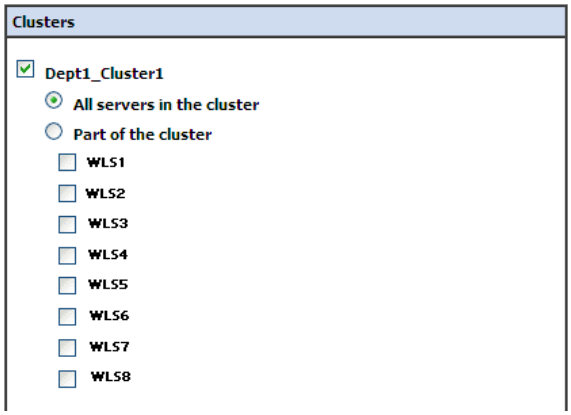

11. Click **Next**.
12. In the **Optional Settings** page, accept the default values, and Click **Next**.
13. Review the configuration settings you have chosen, and select **No, I will review the configuration later** to immediately update the application's configuration after you install it. Then click **Finish** to complete the installation.

8.4 Starting the Web Application

To start the Web application:

1. In the left pane of the Administration Console, select **Deployments**.
2. In the right pane, select the check boxes next to **dizzyworld.ear**.
3. Click **Start** and choose **Servicing all requests** to the Web application.
4. Click **Yes** to confirm that you want to start the selected Web application.

8.5 Testing the Application

After starting the web application, test the application as follows:

1. Add a few items to the shopping cart using the `go shopping` link.
2. Verify that your shopping cart includes the items you selected. If the items are in your shopping cart, the application is working fine.

9

Managing the Topology

This chapter describes some operations that you can perform after you have set up the topology. These operations involve migrating, scaling, and backing up your topology.

This chapter contains the following sections:

9.1 Important Notes Before You Begin

If you are an Oracle Solaris user, read Section 3.1, "Important Notes for Oracle Solaris Users" before you complete the procedures described in this chapter.

9.2 Prerequisites

The following are the prerequisites for managing your enterprise deployment:

- Preconfigure the database and environment, as described in Section 3, "Network, Storage, and Database Preconfiguration".
- Ensure that you have created and configured the Managed Servers in a cluster for `ComputeNode1` and `ComputeNode2`, as described in Section 5, "Configuring Oracle Fusion Middleware".
- Your Node Manager should be configured on `ComputeNode1` and `ComputeNode2`, as described in Section 5.7, "Configuring Java Node Manager".
- Ensure that you have activated the Exalogic-specific features and enhancements in Oracle WebLogic, as described in Chapter 7, "Enabling Exalogic-Specific Enhancements in Oracle WebLogic Server 11g Release 1 (10.3.4)".

- Ensure that your Administration Server is up and running. To start the Administration Server, see Section 5.6, "Starting the Administration Server on ComputeNode1".

9.3 Configuring Server Migration

Server migration is required for applications that have critical data, such as persistent JMS or transaction logs, that needs to be recovered quickly in the event of a failure.

In this enterprise deployment topology, the sample application `dizzyworld.ear` (see, Chapter 8, "Deploying a Sample Web Application to an Oracle WebLogic Cluster") does not leverage JMS or JTA features, so server migration is not required for this sample application. If your applications do require persistent JMS or JTA features, then you may implement server migration.

Managed Servers running on `ComputeNode1` are configured to restart on `ComputeNode2` when a failure occurs. Similarly, Managed Servers running on `ComputeNode2` are configured to restart on `ComputeNode1` if a failure occurs. For this configuration, the Managed Servers running on `ComputeNode1` and `ComputeNode2` listen on specific floating IPs that are failed over by Oracle WebLogic Server Migration. For more information, see Section 3.3.3, "Enterprise Deployment Network Configuration".

Configuring server migration for Oracle Exalogic Managed Servers involves the following steps:

- Setting Up a User and Tablespace for the Server Migration Leasing Table
- Editing the Node Manager's Properties File
- Setting Environment and Superuser Privileges for the wlsifconfig.sh Script
- Configuring Server Migration Targets
- Testing the Server Migration

9.3.1 Prerequisites

The following are the prerequisites for configuring server migration:

- Ensure that the Grid Link Data Sources are created for `ComputeNode1` and `ComputeNode2`, as described in Section 7.6, "Configuring Grid Link Data Source for Dept1_Cluster1".
- Complete the mandatory patch requirements.

  When leveraging WebLogic Server network channels that map to multiple network interfaces with the WebLogic Server migration framework, review the Knowledge Base article (Title: *Oracle Exalogic Elastic Cloud 11g R1 - Known Issues* Doc ID: **1268557.1**) located at My Oracle Support.

  1. Enter the My Oracle Support URL (`https://support.oracle.com/`) in a Web browser.
  2. Click **Sign In** and enter your My Oracle Support username and password.
  3. Search for Doc ID: **1268557.1** in the Search Knowledge Base search box, which is the Document ID of the article related to this issue.

9.3.2 Setting Up a User and Tablespace for the Server Migration Leasing Table

Set up a user and tablespace for the server migration leasing table as follows:

1. Ensure that the database connectivity for your Oracle Exalogic Machine is established, as described in Section 3.5.2, "Connecting to Oracle Database Over Ethernet". Alternatively, if you are connecting your Oracle Exalogic machine to Oracle Exadata Database Machine via InfiniBand, verify the database connectivity and access.
2. Create a tablespace named `leasing`. For example, log on to SQL*Plus as the `sysdba` user and run the following command:

```
SQL> create tablespace leasing
       logging datafile 'DB_HOME/oradata/orcl/leasing.dbf'
       size 32m autoextend on next 32m maxsize 2048m extent management local;
```

> **Note:** Creating a table space is optional.

3. Create a user named `leasing` and assign to it the leasing tablespace.

```
SQL> create user leasing identified by welcome1;

SQL> grant create table to leasing;
SQL> grant create session to leasing;
SQL> alter user leasing default tablespace leasing;
SQL> alter user leasing quota unlimited on LEASING;
```

4. Log in as the `leasing` user, and create the leasing table using the `leasing.ddl` script.
   a. Copy the `leasing.ddl` file located in the `/u01/app/FMW_Product1/Oracle/Middleware/wlserver_10.3/server/db2/oracle/920` directory to your database node.
   b. Connect to the database as the `leasing` user.
   c. Run the `leasing.ddl` script in SQL*Plus.

```
SQL> @copy_location/leasing.ddl;
```

9.3.3 Creating a GridLink Source Using the Oracle WebLogic Administration Console

Ensure that you have created the GridLink Data Source, as described in Section 7.6.2, "Creating a GridLink Data Source on Dept1_Cluster1".

9.3.4 Editing the Node Manager's Properties File

You must complete this task for the node managers on both `ComputeNode1` and `ComputeNode2` where server migration is being configured. The `nodemanager.properties` file is located in the `/u01/Dept_1/admin/el01cn01/nodemanager` directory on `ComputeNode1` and `/u01/Dept_1/admin/el01cn02/nodemanager` on `ComputeNode2`:

> **Note:** Ensure that you edit the `nodemanager.properties` files of both `ComputeNode1` and `ComputeNode2`.

```
bond0=10.0.0.1-10.0.0.17,NetMask=255.255.255.224
bond1=10.1.0.1-10.1.0.17,NetMask=255.255.255.224
UseMACBroadcast=true
```

Verify in Node Manager's output (shell where Node Manager is started) that these properties are being used, or problems may arise during migration. You should see something like this in the Node Manager's output:

```
bond0=10.0.0.1-10.0.0.17,NetMask=255.255.255.224
bond1=10.1.0.1-10.1.0.17,NetMask=255.255.255.224
UseMACBroadcast=true
```

For more information, see the "Reviewing nodemanager.properties" section in the *Oracle Fusion Middleware Node Manager Administrator's Guide for Oracle WebLogic Server*.

> **Note:** The steps below are not required if the server properties (start properties) have been properly set and the Node Manager can start the servers remotely.

1. Set the following property in the `nodemanager.properties` file.
   - `StartScriptEnabled`

     Set this property to `true`. This is required for the shiphome to enable the Node Manager to start the managed servers.
2. Start the Node Manager on `ComputeNode1` and `ComputeNode2` by running the `startNodeManager.sh` script. This scripts are located at:
   - In `ComputeNode1`: `/u01/Dept1/admin/el01cn01/nodemanager` directory
   - In `ComputeNode2`: `/u01/Dept1/admin/el01cn02/nodemanager` directory

> **Note:** When running Node Manager from a shared storage installation, multiple nodes are started using the same `nodemanager.properties` file. However, each node may require different `NetMask` or `Interface` properties. In this case, specify individual parameters on a per-node basis using environment variables. For example, to use a different interface in `ComputeNode`*n*, use the `Interface` environment variable as follows: `ComputeNoden> export JAVA_OPTIONS=-DInterface=bond0` and start Node Manager after the variable has been set in the shell.

9.3.5 Setting Environment and Superuser Privileges for the wlsifconfig.sh Script

Set environment and superuser privileges for the `wlsifconfig.sh` script:

1. Ensure that your `PATH` environment variable includes these files:

Table 9–1 Files Required for the PATH Environment Variable

File	Located in this directory
`wlsifconfig.sh`	`/u01/app/FMW_Product1/Oracle/Middleware/wlserver_10.3/common/bin`
`wlscontrol.sh`	`/u01/app/FMW_Product1/Oracle/Middleware/wlserver_10.3/common/bin`

Table 9–1 (Cont.) Files Required for the PATH Environment Variable

File	Located in this directory
nodemanager.domains	/u01/Dept_1/admin/el01cn01/nodemanager on ComputeNode1 and /u01/Dept_1/admin/el01cn02/nodemanager on ComputeNode2

2. Run `wlsifconfig.sh -listif bond0` script and verify that your network interface and the netmask (for example, `255.255.255.192` for the `Dept_1` subnet used in the example configuration) are correct.
3. Grant sudo configuration for the `wlsifconfig.sh script`.

Note: Ensure that you run `sudo /sbin/ifconfig` or `sudo /sbin/arping`. Running this command will disable any password input prompt.

- Configure `sudo` to work without a password prompt.
- For security reasons, sudo should be restricted to the subset of commands required to run the `wlsifconfig.sh` script. For example, to set the environment and superuser privileges for the `wlsifconfig.sh` script, complete these steps:
  a. Grant sudo privilege to the WebLogic user (`weblogic`) with no password restriction, and grant execute privileges on the */sbin/ifconfig* and */sbin/arping* binaries.
  b. Make sure that the script is executable by the WebLogic user (weblogic). The following is an example of an entry inside */etc/sudoers* granting sudo execution privilege for `weblogic` and also over `ifconfig` and `arping`:

```
weblogic ALL=NOPASSWD: /sbin/ifconfig,/sbin/arping
```

Note: Contact your system administrator for the sudo and system rights as appropriate to this step.

9.3.6 Configuring Server Migration Targets

You must configure server migration targets. To do so, complete the following:

- Configuring Server Migration Targets for Dept1_Cluster1
- Configuring Server Migration Targets for Managed Servers Running onComputeNode1

9.3.6.1 Configuring Server Migration Targets for Dept1_Cluster1

You must configure server migration targets for the cluster (`Dept1_Cluster1`). Configuring Cluster Migration sets the `DataSourceForAutomaticMigration` property to true. Follow the steps below to configure cluster migration in a cluster:

1. Log in to the Oracle WebLogic Server Administration Console, using the following URL:

```
http://ADMINVHN1:7001/console
```

2. If you have not already done so, in the Change Center of the Administration Console, click **Lock & Edit**.
3. In the left pane of the Console, expand **Environment** and then select **Clusters**.

   The **Summary of Clusters** page is displayed.
4. Click **Dept1_Cluster1** for which you want to configure migration in the Name column of the table.

   The **Settings for Dept1_Cluster1** page is displayed.
5. Click the **Migration** tab.
6. Enter the following details:
   - For the **Candidate Machines For Migratable Servers**, select **ComputeNode2** under **Available**, and then click the right arrow.
   - For **Migration Basis**, select **Database**.
   - For **Data Source For Automatic Migration**, select **gridlink**. This is the data source you created in Section 7.6.2, "Creating a GridLink Data Source on Dept1_Cluster1".
7. Click **Save**.
8. Click **Activate Changes**.

9.3.6.2 Configuring Server Migration Targets for Managed Servers Running onComputeNode1

Set the Managed Servers on `ComputeNode1` for server migration as follows:

1. If you have not already done so, in the Change Center of the Administration Console, click **Lock & Edit**.
2. In the left pane of the Console, expand **Environment** and then select **Servers**.
3. Select **WLS1**.

   The **Settings for WLS1** is displayed.
4. Click the **Migration** tab.
5. In the **Migration Configuration** page, enter the following details:
   a. Select **Automatic Server Migration Enabled**. This enables the Node Manager to start a failed server on the target node automatically.
   b. For **Candidate Machines**, select **ComputeNode2** under **Available** and click the right arrow.
   c. For **JMS Service Candidate** Servers, select all of the Managed Servers (Table 5–2) in `ComputeNode2` and click the right arrow.
   d. Select **Automatic JTA Migration Enabled**. This enables the automatic migration of the JTA Transaction Recovery System on this server.
   e. For **JTA Candidate Servers**, select all of the Managed Servers (Table 5–2) in `ComputeNode2` and click the right arrow.
6. Click **Save**.
7. Click **Activate Changes**.
8. Restart the Administration Server and the servers for which server migration has been configured.

To restart the Administration Server, use the procedure in Section 5.8, "Restarting the Administration Server on ComputeNode1."

Tip: Click **Customize this table** in the Summary of Servers page, move Current Machine from the Available Window to the Chosen window to view the machine on which the server is running. This will be different from the configuration if the server gets migrated automatically.

9.3.7 Testing the Server Migration

You must test the server migration. To verify that Server Migration is working properly, follow these steps:

Notes:

- The migratable IP address should not be present on the interface of any of the candidate machines before the migratable server is started.
- Ensure that a minimum of two Managed Servers in the cluster are up and running before you test the server migration.

Testing from ComputeNode1:

1. Stop the `ComputeNode1` Managed Server.

   To do this, run this command:

   ```
   ComputeNode1> kill -9 <pid>
   ```

   pid specifies the process ID of the Managed Server (`WLS1`). You can identify the pid in the node by running this command:

   ```
   ComputeNode1> ps -ef | grep WLS1
   ```

2. Watch the Node Manager console: you should see a message indicating that `ComputeNode1`'s floating IP has been disabled.
3. Wait for the Node Manager to try a second restart of `ComputeNode1`. Node Manager waits for a fence period of 30 seconds before trying this restart.
4. Once Node Manager restarts the server, stop it again. Now Node Manager should log a message indicating that the server will not be restarted again locally.

Testing from ComputeNode2:

Watch the local Node Manager console. After 30 seconds since the last try to restart Node Manager on `ComputeNode1`, Node Manager on `ComputeNode2` should prompt that the IP for `ComputeNode1` is being brought up and that the server is being restarted in this compute node.

Verifying from the Administration Console

You can also verify server migration from the Administration Console as follows:

1. Log in to the Administration Console.
2. Click on **Domain** on the left pane.
3. Click the **Monitoring** tab and then the **Migration** subtab.

The Migration Status table provides information on the status of the migration.

9.4 Connecting Two Subnets Used by Different Departments

By completing the configuration procedures described in Chapter 5, "Configuring Oracle Fusion Middleware", you have set up and configured the environment for `Dept_1`, which uses `ComputeNode1` and `ComputeNode2`. Similarly, you can configure the environment for another department, such as `Dept_2`, which uses `ComputeNode3` and `ComputeNode4`.

If you wish to isolate the application deployment and environment for `Dept_1` from that of `Dept_2`, then you must create separate IP subnets for both `Dept_1` and `Dept_2` over the default IP over InfiniBand (IPoIB) link. For more information about creating such subnets, see Section 3.3.3.5.1, "Application Isolation by IP Subnetting over IPoIB".

In some scenarios, the `Dept_1` application may require communication with the `Dept_2` application. To enable the `Dept_1` application (deployed on `ComputeNode1` and `ComputeNode2`) to communicate with the `Dept_2` application (deployed on `ComputeNode3` and `ComputeNode4`), you must set up IP aliasing for the two subnets to access each other. To set up this IP aliasing, see Section 3.3.3.5.2, "Allowing a Compute Node to Access Two Different Subnets Simultaneously".

9.5 Scaling Out the Topology - Adding Managed Servers to New Compute Nodes

When you scale out the topology, you add a new Managed Server to new compute nodes, such as `ComputeNode3`. In this example procedure, `WLS9` is created on `ComputeNode3` and added to `Dept1_Cluster1`.

9.5.1 Prerequisites

Before performing the steps in this section, verify that your environment meets the following requirements:

- There must be existing compute nodes, such as `ComputeNode1` and `ComputeNode2`, running Managed Servers configured with Oracle WebLogic Server within the topology.
- The new compute node, such as `ComputeNode3`, can access the existing home directories for Oracle WebLogic Server whose binaries are installed in a separate share on the Sun ZFS Storage 7320 appliance. For more information, see Section 3.4, "Shared Storage and Recommended Project and Share Structure".
- In the Exalogic environment, ORACLE_HOME and WL_HOME directories are shared by multiple servers in different compute nodes. Therefore, Oracle recommends that you keep the Oracle Inventory and Middleware home list in those nodes updated for consistency in the installations and application of patches. To update the oraInventory in a compute node and "attach" an installation in the shared file system on Sun ZFS Storage 7320 appliance to the inventory, use the *ORACLE_HOME/oui/bin/attachHome.sh* script. To update the Middleware home list to add or remove a WL_HOME, edit the *<user_home>/bea/beahomelist* file. See the steps below.

9.5.2 Adding Managed Servers to ComputeNode3

You must add the Managed Servers on `ComputeNode3` as follows:

- Mounting Existing Oracle Fusion Middleware Home and Domain on ComputeNode3
- Propagating Domain Configuration from ComputeNode1 to ComputeNode3 Using pack and unpack Utilities
- Setting Up Java-Based Node Manager on ComputeNode3
- Setting Environment and Superuser Privileges for the wlsifconfig.sh Script
- Creating and Configuring a Machine
- Creating a Managed Server on ComputeNode3
- Assigning the Managed Server to the New Machine
- Configuring Network Channels for Managed Servers on ComputeNode3
- Configuring Persistent Store
- Starting Node Manager on ComputeNode3
- Starting Managed Servers on ComputeNode3
- Configuring Server Migration Targets
- Testing Server Migration

9.5.2.1 Mounting Existing Oracle Fusion Middleware Home and Domain on ComputeNode3

On `ComputeNode3`, mount the existing Oracle Fusion Middleware Home, which should include the Oracle WebLogic Server installation (Located at `/u01/app/FMW_Product1/Oracle/Middleware`) and the domain directory (Located at `/u01/Dept_1/domains/el01cn03`), and ensure that the new compute node (`ComputeNode3`) has access to this directory, just like the rest of the compute nodes in the domain (`ComputeNode1` and `ComputeNode2`). You must complete the following:

- Creating Mount Points on ComputeNode3
- Editing the /etc/fstab File (Linux) or /etc/vfstab (Solaris)
- Mounting the Volumes

9.5.2.1.1 Creating Mount Points on ComputeNode3 On the command line, run the following commands on `ComputeNode3` to create the necessary mount points:

```
# mkdir -p /u01/common/patches
# mkdir -p /u01/common/general
# mkdir -p /u01/FMW_Product1/Oracle/wlserver_10.3
# mkdir -p /u01/FMW_Product1/webtier_1115
# mkdir -p /u01/Dept_1/domains/el01cn03
# mkdir -p /u01/el01cn03/dumps
# mkdir -p /u01/el01cn03/general
```

where `el01cn03` is the host name assigned to `ComputeNode3`.

9.5.2.1.2 Editing the /etc/fstab File (Linux) or /etc/vfstab (Solaris) After creating the mount points, you must add entries for the mount points to the `/etc/fstab` (Linux) or `/etc/vfstab` (Solaris) file on `ComputeNode3`.

On `ComputeNode3`, log in as a `root` user and add the following entries to the `/etc/fstab` (Linux) or `/etc/vfstab` (Solaris)file in a text editor, such as `vi`:

Oracle Linux

- el01sn01-priv:/export/common/general /u01/common/general nfs4 rw,bg,hard,nointr,rsize=131072,wsize=131072
- el01sn01-priv:/export/common/patches /u01/common/patches nfs4 rw,bg,hard,nointr,rsize=131072,wsize=131072
- el01sn01-priv:/export/el01cn03/dumps /u01/el01cn03/dumps nfs4 rw,bg,hard,nointr,rsize=131072,wsize=131072
- el01sn01-priv:/export/el01cn03/general /u01/el01cn03/general nfs4 rw,bg,hard,nointr,rsize=131072,wsize=131072
- el01sn01-priv:/export/Dept_1/domains /u01/Dept_1/domains nfs4 rw,bg,hard,nointr,rsize=131072,wsize=131072
- el01sn01-priv:/export/Dept_1/jmsjta /u01/Dept_1/jmsjta nfs4 rw,bg,hard,nointr,rsize=135268,wsize=135168
- el01sn01-priv:/export/Dept_1/admin /u01/Dept_1/admin nfs4 rw,bg,hard,nointr,rsize=131072,wsize=131072
- el01sn01-priv:/export/FMW_Product1/wlserver_1034 /u01/FMW_Product1/wlserver_1034 nfs4 rw,bg,hard,nointr,rsize=131072,wsize=131072
- el01sn01-priv:/export/FMW_Product1/webtier_1115 /u01/FMW_Product1/webtier_1115 nfs4 rw,bg,hard,nointr,rsize=131072,wsize=131072

Oracle Solaris

- el01sn01-priv:/export/common/general - /u01/common/general nfs - yes rw,bg,hard,nointr,rsize=131072,wsize=131072,vers=4
- el01sn01-priv:/export/common/patches - /u01/common/patches nfs - yes rw,bg,hard,nointr,rsize=131072,wsize=131072,vers=4
- el01sn01-priv:/export/el01cn03/dumps - /u01/el01cn03/dumps nfs - yes rw,bg,hard,nointr,rsize=131072,wsize=131072,vers=4
- el01sn01-priv:/export/el01cn03/general - /u01/el01cn03/general nfs - yes rw,bg,hard,nointr,rsize=131072,wsize=131072,vers=4
- el01sn01-priv:/export/Dept_1/domains - /u01/Dept_1/domains nfs - yes rw,bg,hard,nointr,rsize=131072,wsize=131072,vers=4
- el01sn01-priv:/export/Dept_1/jmsjta - /u01/Dept_1/jmsjta nfs - yes rw,bg,hard,nointr,rsize=135268,wsize=135168,vers=4
- el01sn01-priv:/export/Dept_1/admin - /u01/Dept_1/admin nfs - yes rw,bg,hard,nointr,rsize=131072,wsize=131072,vers=4
- el01sn01-priv:/export/FMW_Product1/wlserver_1034 - /u01/FMW_Product1/wlserver_1034 nfs - yes rw,bg,hard,nointr,rsize=131072,wsize=131072,vers=4
- el01sn01-priv:/export/FMW_Product1/webtier_1115 - /u01/FMW_Product1/webtier_1115 nfs - yes rw,bg,hard,nointr,rsize=131072,wsize=131072,vers=4

Note: In the above entries, `el01sn01-priv` is used as the example host name of the Sun ZFS Storage 7320 appliance. You can also use the IPoIB IP address assigned to the storage appliance.

Save the file and exit.

On the command line, run the following commands as a `root` user on `ComputeNode1` and `ComputeNode2` to create the necessary mount points:

```
# mkdir -p /u01/Dept_1/admin/el01cn03/nodemanager
# mkdir -p /u01/Dept_1/jmsjta/base_domain/Dept1_Cluster1/jms
# mkdir -p /u01/Dept_1/jmsjta/base_domain/Dept1_Cluster1/tlogs
```

9.5.2.1.3 Mounting the Volumes

To mount the volumes, complete the following steps:

1. On `ComputeNode3`, ensure that the mount entries are added to the `/etc/fstab` file correctly.
2. Run the `mount -a` command on `ComputeNode3` to mount the volumes.

9.5.2.2 Propagating Domain Configuration from ComputeNode1 to ComputeNode3 Using pack and unpack Utilities

You have created the domain (`base_domain`) on `ComputeNode1`. You must propagate the domain configuration to `ComputeNode3` as follows:

1. Run the `pack` command on `ComputeNode1` to create a template pack using the following commands:

```
ComputeNode1> cd /u01/app/FMW_Product1/Oracle/Middleware/wlserver_
10.3/common/bin

ComputeNode1> ./pack.sh -managed=true -domain=/u01/Dept_
1/domains/el01cn01/base_domain -template=basedomaintemplate.jar -template_
name=basedomain_template
```

2. Run the `unpack` command on `ComputeNode3` to unpack the template.

```
ComputeNode3> cd /u01/app/FMW_Product1/Oracle/Middleware/wlserver_
10.3/common/bin

ComputeNode3> ./unpack.sh -/u01/Dept_1/domains/el01cn03/base_domain
-template=basedomaintemplate.jar
```

9.5.2.3 Setting Up Java-Based Node Manager on ComputeNode3

Complete the procedure described in Section 5.7, "Configuring Java Node Manager" and ensure that you make the following changes:

- New Node Manager directory: `u01/Dept_1/admin/el01cn03/nodemanager`
- Listen Address= `192.168.10.3`

Note: This IP address is the `Bond0` IP address of `ComputeNode3`.

- Domain Home= `/u01/Dept_1/domains/el01cn03/base_domain`

- DomainsFile= /u01/Dept_1/admin/el01cn03/nodemanager/nodemanager.domains
- LogFile= /u01/Dept_1/admin/el01cn03/nodemanager/nodemanager.log

In addition, complete the following steps on ComputeNode3:

1. Start WLST as follows:

```
ComputeNode1> cd /u01/app/FMW_Product1/Oracle/Middleware/wlserver_
10.3/common/bin
ComputeNode1> ./wlst.sh
```

2. Use the connect command to connect WLST to a WebLogic Server instance, as in the following example:

```
wls:/offline> connect('username','password','t3://ADMINVHN1:7001')
```

3. Once you are in the WLST shell, run nmEnroll using the following syntax:

```
nmEnroll([domainDir], [nmHome])
```

   For example,

```
nmEnroll('/u01/Dept_1/domains/el01cn03/base_domain', '/u01/Dept_
1/admin/el01cn03/nodemanager')
```

   Running nmEnroll ensures that the correct Node Manager user and password token are supplied to each Managed Server. Once these are available for each Managed Server, you can use nmConnect in a production environment.

4. Disconnect WLST from the WebLogic Server instance by entering disconnect(), and exit by entering exit() to exit the WLST shell.

9.5.2.4 Setting Environment and Superuser Privileges for the wlsifconfig.sh Script

Set environment and superuser privileges for the wlsifconfig.sh script:

1. Ensure that your PATH environment variable includes these files:

Table 9–2 Files Required for the PATH Environment Variable

File	Located in this directory
wlsifconfig.sh	/u01/app/FMW_Product1/Oracle/Middleware/wlserver_10.3/common/bin
wlscontrol.sh	/u01/app/FMW_Product1/Oracle/Middleware/wlserver_10.3/common/bin
nodemanager.domains	/u01/Dept_1/admin/el01cn03/nodemanager on ComputeNode3

2. Run wlsifconfig.sh -listif bond0 script and verify that your network interface and the netmask (for example, 255.255.255.192 for the Dept_1 subnet used in the example configuration) are correct.
3. Grant sudo configuration for the wlsifconfig.sh script.

Note: Ensure that you run `sudo /sbin/ifconfig` or `sudo /sbin/arping`. Running this command will disable any password input prompt.

- Configure `sudo` to work without a password prompt.
- For security reasons, sudo should be restricted to the subset of commands required to run the `wlsifconfig.sh` script. For example, to set the environment and superuser privileges for the `wlsifconfig.sh` script, complete these steps:
  - **a.** Grant sudo privilege to the WebLogic user (`weblogic`) with no password restriction, and grant execute privileges on the */sbin/ifconfig* and */sbin/arping* binaries.
  - **b.** Make sure that the script is executable by the WebLogic user (weblogic). The following is an example of an entry inside */etc/sudoers* granting sudo execution privilege for `weblogic` and also over `ifconfig` and `arping`:

    ```
    weblogic ALL=NOPASSWD: /sbin/ifconfig,/sbin/arping
    ```

Note: Contact your system administrator for the sudo and system rights as appropriate to this step.

9.5.2.5 Creating and Configuring a Machine

Create a new machine for `ComputeNode3` that will be used, and add the machine to the domain. Complete the following steps:

1. Log in to the Oracle WebLogic Administration Console.
2. If you have not already done so, in the Change Center of the Administration Console, click **Lock & Edit**.
3. In the left pane of the Console, expand **Environment** and select **Machines**.

   The **Summary of Machines** page is displayed.
4. Click **New**.

   The **Create a New Machine** page is displayed.
5. Enter the following details:
   - Enter **ComputeNode3** as the name for the new machine in the **Name** field.
   - Select **UNIX** from the drop-down list in the **Machine OS** field.
   - Click **Next**.

   The **Node Manager Properties** page is displayed.
6. Enter the following details:
   - Select **Plain** from the drop-down list in the **Type** field.
   - **Listen Address**: `192.168.10.3`

Note: This address is the example `BOND0` IP address of `ComputeNode3`.

- **Listen Port**: `5556`
- **Node Manager Home**: `/u01/Dept_1/admin/el01cn03/nodemanager`

7. Click **Finish**.

   The new machine is displayed in the Machines table.

9.5.2.6 Creating a Managed Server on ComputeNode3

You can create a Managed Server on `ComputeNode3` in an existing domain (base_domain) which is shared by `ComputeNode1` and `ComputeNode2`. Complete the following:

1. If you have not already done so, in the Change Center of the Administration Console, click **Lock & Edit**.
2. In the left pane of the Console, expand **Environment**, and then **Servers**.

   The **Summary of Servers** page is displayed.
3. In the **Servers** table, click **New**.

   The **Create a New Server** page is displayed.
4. Enter the following details:
   - Enter **WLS9** as the name of the server in the **Name** field.
   - **Listen Address**: `10.0.0.9`

     Note: This address is the floating IP address assigned to the new Managed Server. This address uses the `BOND0` interface. Ensure that the address is configured before you start the new Managed Server.

   - **Listen Port**: `7003`
   - Check **Yes, make this server a member of an existing cluster**, and then select **Dept1_Cluster1**.

     Click **Next**.
5. Review the configuration options you have chosen, and click **Finish**.
6. Click **Activate Changes**.

9.5.2.7 Assigning the Managed Server to the New Machine

You must associate the Managed Servers you created to the new machine (`ComputeNode3`). Complete the following steps:

1. If you have not already done so, in the Change Center of the Administration Console, click **Lock & Edit**.
2. In the left pane of the console, expand **Environment**, and then **Servers**.

   The **Summary of Servers** page is displayed.
3. Select **WLS9**.

   The **Settings for WLS9** page is displayed.
4. Select **Configuration**, and then **General**.
5. In the Machine field, select **ComputeNode3**.

6. Click **Save**.
7. To activate these changes, in the Change Center of the Administration Console, click **Activate Changes**.

9.5.2.8 Configuring Network Channels for Managed Servers on ComputeNode3

For the Managed Server on `ComputeNode3` (`WLS9`), you can create the following network channels:

- HTTP Client Channel
- T3 Client Channel

9.5.2.8.1 HTTP Client Channel Create the network channel for the Managed Server on `ComputeNode3`, by following the instructions described in Section 5.12.3.1, "HTTP Client Channel" and enter the required properties as described in Table 9–3.

Table 9–3 Network Channels Properties

Managed Server	Name	Protocol	Listen Address	Listen Port
WLS9	HTTPClient9	http	10.1.0.9	7003

9.5.2.8.2 T3 Client Channel Create the network channel for the Managed Server on `ComputeNode3`, by following the instructions described in Section 5.12.3.2, "T3 Client Channel" and enter the required properties as described in Table 9–4.

Table 9–4 Network Channels Properties

Managed Server	Name	Protocol	Listen Address	Listen Port
WLS9	T3ClientChannel9	t3	10.1.0.9	7003

Note: These IPs use the `Bond1` interface.

9.5.2.9 Configuring Persistent Store

Configure the persistent store for the new server. This should be a location visible from other compute nodes, as recommended in Section 3.4, "Shared Storage and Recommended Project and Share Structure."

From the Administration Console, select the **Server_name**, and then the **Services** tab. Under **Default Store**, in **Directory**, enter the path to the folder where you want the default persistent store to store its data files.

9.5.2.10 Starting Node Manager on ComputeNode3

Start Node Manager on the new compute node (`ComputeNode3`). To start Node Manager, use the installation in shared storage from the existing compute nodes (`ComputeNode1` or `ComputeNode2`), and start Node Manager by passing the host name of the new node as a parameter as follows:

```
ComputeNode3> /u01/Dept_1/admin/el01cn03/nodemanager/startNodeManager.sh
```

9.5.2.11 Starting Managed Servers on ComputeNode3

Start and test the new Managed Server from the Oracle WebLogic Server Administration Console.

1. Shut down all the existing Managed Servers in the cluster.
2. Ensure that the newly created Managed Servers `WLS9` is running.

9.5.2.12 Configuring Server Migration Targets

You must configure server migration targets. To do so, complete the following:

- Configuring Server Migration Targets for Dept1_Cluster1
- Configuring Server Migration Targets for Managed Servers Running onComputeNode1

9.5.2.12.1 Configuring Server Migration Targets for Dept1_Cluster1 You must configure server migration targets for the Dept1_Cluster1 cluster. Configuring Cluster Migration sets the `DataSourceForAutomaticMigration` property to true. Follow the steps below to configure cluster migration in a cluster:

1. Log in to the Oracle WebLogic Server Administration Console, using the following URL:

   ```
   http://ADMINVHN1:7001/console
   ```

2. If you have not already done so, in the Change Center of the Administration Console, click **Lock & Edit**.
3. In the left pane of the Console, expand **Environment** and then select **Clusters**.

   The **Summary of Clusters** page is displayed.
4. Click **Dept1_Cluster1** for which you want to configure migration in the Name column of the table.

   The **Settings for Dept1_Cluster1** page is displayed.
5. Click the **Migration** tab.
6. Enter the following details:
   - For the **Candidate Machines For Migratable Servers**, select **ComputeNode2** under **Available**, and then click the right arrow.
   - For **Migration Basis**, select **Database**.
   - For **Data Source For Automatic Migration**, select **gridlink**. This is the data source you created in Section 7.6.2, "Creating a GridLink Data Source on Dept1_Cluster1".
7. Click **Save**.
8. Click **Activate Changes**.

9.5.2.12.2 Configuring Server Migration Targets for the Managed Server Running on ComputeNode3 Set the Managed Servers on `ComputeNode3` for server migration as follows:

1. If you have not already done so, in the Change Center of the Administration Console, click **Lock & Edit**.
2. In the left pane of the Console, expand **Environment** and then select **Servers**.
3. Select **WLS9**.

   The **Settings for WLS9** is displayed.
4. Click the **Migration** tab.

5. In the **Migration Configuration** page, enter the following details:
   a. Select **Automatic Server Migration Enabled**. This enables the Node Manager to start a failed server on the target node automatically.
   b. For **Candidate Machines**, select **ComputeNode3** under **Available** and click the right arrow.
   c. For **JMS Service Candidate** Servers, select all of the Managed Servers (Table 5–2) in `ComputeNode2` and click the right arrow.
   d. Select **Automatic JTA Migration Enabled**. This enables the automatic migration of the JTA Transaction Recovery System on this server.
   e. For **JTA Candidate Servers**, select all of the Managed Servers (Table 5–2) in `ComputeNode2` and click the right arrow.
6. Click **Save**.
7. Click **Activate Changes**.
8. Restart the Administration Server and the servers for which server migration has been configured.

   To restart the Administration Server, use the procedure in Section 5.8, "Restarting the Administration Server on ComputeNode1."

   Tip: Click **Customize this table** in the Summary of Servers page, move Current Machine from the Available Window to the Chosen window to view the machine on which the server is running. This will be different from the configuration if the server gets migrated automatically.

9.5.2.13 Testing Server Migration

Test server migration for this new server. Follow these steps from the node where you added the new server:

1. Abruptly stop the `WLS9` Managed Server by running `kill -9 <pid>` on the PID of the Managed Server. You can identify the PID of the node using `ps -ef | grep WLSn`.
2. In the Node Manager Console, you should see a message indicating that WLS9's floating IP has been disabled.
3. Wait for the Node Manager to try a second restart of WLS*n*. Node Manager waits for a fence period of 30 seconds before trying this restart.
4. Once Node Manager restarts the server, stop it again. Now Node Manager should log a message indicating that the server will not be restarted again locally.

9.6 Scaling Down the Topology: Deleting Managed Servers

To scale down the topology by deleting the new Managed Servers on ComputeNode3, complete the following steps:

- Deleting a Managed Server
- Deleting the Machine

9.6.1 Deleting a Managed Server

When you delete a Managed Server, WebLogic Server removes its associated configuration data from the domain's configuration file (`config.xml`). All of the configuration data for the server will be deleted. For example, any network channels that you created for the server are deleted, but applications and EJBs that are deployed on the server will not be deleted.

> **Note:** Ensure that you have stopped the Managed Server, before you delete it.

To delete a Managed Server such as `WLS9`:

1. Log in to the Oracle WebLogic Administration Console.
2. If you have not already done so, click **Lock & Edit** in the Change Center.
3. In the left pane of the Console, select **Environment**, and then **Servers**.

   The **Summary of Servers** page is displayed.
4. Select the check box next to **WLS9** in the Names column of the table and click **Delete**.
5. Confirm your deletion request.
6. Click **Activate Changes**.

9.6.2 Deleting the Machine

You must delete the machine (`ComputeNode3`) on `ComputeNode3` by completing the following steps:

1. Log in to the Oracle WebLogic Administration Console.
2. If you have not already done so, in the Change Center of the Administration Console, click **Lock & Edit**.
3. In the left pane of the Console, expand **Environment** and select **Machines**.

   The **Summary of Machines** page is displayed.
4. Select the check box next to **ComputeNode3** and click **Delete**.

   A dialog displays asking you to confirm your deletion request.
5. Click **Yes** to delete the machine.
6. To activate these changes, in the Change Center of the Administration Console, click **Activate Changes**.

9.7 Performing Backups and Recoveries

This section describes backup and recovery recommendations for Oracle Exalogic users.

It contains the following topics:

- Important Artifacts to Back Up
- boot.properties File for Restored Managed Servers

Important Artifacts to Back Up

Table 9–5 lists the WebLogic artifacts that you should back up frequently in the Oracle Exalogic enterprise deployment.

Table 9–5 WebLogic Artifacts to Back Up in the Oracle Exalogic Enterprise Deployment

Type	Location
JMS persistent messages and JTA tlogs	`/u01/Dept_1/jmsjta/base_domain/Dept1_Cluster1/jms` and `/u01/Dept_1/jmsjta/base_domain/Dept1_Cluster1/tlogs`
Read-write home directories, such as OS syslogs and crash dump	`/u01/el01cn01/dumps` and `/u01/el01cn01/general`
Read-only home directories, such as ORACLE_HOME, OS patches, and so on.	User-defined share on the Sun ZFS Storage 7320 appliance
Domain Home directories of all Exalogic compute nodes. Each compute node has its own Domain Home directory. Backing up the configuration data is important to preserve the domain configuration changes.	User-defined share on the Sun ZFS Storage 7320 appliance

For more information about backup and recovery, refer to the following sections, in the *Oracle Fusion Middleware Administrator's Guide*.

- Backing Up Your Environment
- Recovering Your Environment

boot.properties File for Restored Managed Servers

If a Managed Server is restored from a backup, you must re-create the `boot.properties` file located in `/u01/Dept_1/domains/el01cn01/base_domain/servers/AdminServer/security` for the restored Managed Server on its compute node.

To do this, perform these steps:

1. Create the above directory if it does not already exist.
2. In a text editor, create a file called *boot.properties* in the `security` directory created in the previous step, and enter the following lines in the file:

   ```
   username=<admin_user>
   password=<password>
   ```

3. Start the Managed Server.

9.8 Patching Oracle Software and Updating Firmware in Oracle Exalogic Environment

This section provides recommendations and considerations for patching Oracle software and updating firmware in the Oracle Exalogic environment.

It contains the following sections:

- Section 9.8.1, "Oracle Linux"

- Section 9.8.2, "Oracle Solaris"
- Section 9.8.3, "Oracle WebLogic Server"
- Section 9.8.4, "Patching Software or Updating Firmware for Exalogic Machine Hardware Components"

9.8.1 Oracle Linux

If you are an Oracle Linux user, Oracle Linux is pre-installed on each of the Exalogic compute nodes. Any operating system patches reside on the Sun ZFS Storage 7320 appliance, which is used by all compute nodes.

Oracle recommends that you patch the Oracle Linux operating system installed on compute nodes simultaneously. This practice helps you maintain the operating system environment. Therefore, you should patch the operating system at the Exalogic machine level - all compute nodes at once.

9.8.2 Oracle Solaris

If you are updating Oracle Solaris installed on your Exalogic compute nodes, Oracle recommends that you update the Oracle Solaris operating system installed on compute nodes simultaneously. This practice helps you maintain the operating system environment. Therefore, you should update the operating system at the Exalogic machine level - all compute nodes at once.

Updates to Oracle Solaris can be downloaded from the support repository, which is available at the following URL:

`http://pkg.oracle.com/solaris/release`

This URL is restricted. It can only be accessed by the `IPS pkg` commands. Ensure that you register and obtain the key and certificate from `https://pkg-register.oracle.com` before you download any Oracle Solaris 11 Express Support Repository Updates (SRUs).

For more information about the support repository and SRUs (support repository updates), see "Support FAQ" on the Oracle Solaris 11 Express Overview" and "Support Package Repositories Explained".

9.8.3 Oracle WebLogic Server

Oracle Exalogic Machine uses the Sun ZFS Storage 7320 appliance that allows all Oracle WebLogic instances in the Oracle Exalogic system, including instances running in different Oracle WebLogic Server domains, to share the same Oracle WebLogic Server installation.

Topologies using shared installations across Oracle WebLogic Server domains and physical servers offer some advantages over topologies using dedicated installations per domain or physical server. Shared installation topologies create fewer sets of product binaries to be managed, simplify the mapping of Oracle WebLogic Server instances to the installation being used, and enforce maximum consistency of Oracle WebLogic Server versions and maintenance levels in the Oracle Exalogic system. In some environments, shared installation topologies may result in management efficiencies.

However, in some scenarios, you may require multiple Oracle WebLogic Server installations within the Oracle Exalogic system, each dedicated to specific WebLogic Server domains or to compute nodes. Topologies with multiple dedicated installations

provide more management flexibility, particularly when maintenance considerations are important.

Applications running in different WebLogic Server domains may have different maintenance requirements. The frequency of their updates may vary, and the update requirements may affect different functional areas of the WebLogic Server product, resulting in diverse patch requirements. They may also host applications from different departments or business units within the same organization, which require that their applications and systems, including the Oracle WebLogic Server products being used in those applications, are isolated from other applications to minimize cross-application dependencies. Therefore, Oracle recommends that you evaluate your specific WebLogic Server maintenance requirements when determining the installation topology that will be used within the Oracle Exalogic system.

One-Off Patches

One-off patches are provided to address specific functional issues within the WebLogic Server product. One-off patches are generally applied only when required to resolve problems observed in the user's environment. While WebLogic Server patches can be applied on a per-domain basis (or on a more fine-grained basis), Oracle recommends that one-off patches be applied on an installation-wide basis. One-off patches applied to a WebLogic Server installation using this recommended practice affect all domains and servers sharing that installation.

Maintenance Release

Maintenance releases are applied on an installation-wide basis, and once applied, will affect all domains and servers sharing that installation. Oracle recommends that you create a unique Oracle WebLogic Server installation for each set of domains and compute nodes that must be maintained independently of their peers in the Oracle Exalogic system.

Maintenance Type to Evaluate

You should evaluate your specific requirements for maintaining domains and compute nodes within the Oracle Exalogic system and how to group (or isolate) domains and compute nodes from a maintenance perspective.

For example, you can group domains and compute nodes based on the following:

- Departments or business units they support
- Required service levels
- Current and future requirements for isolating domains
- Historical practice

After you arrive at a logical group of domains and compute nodes, you can set up an Oracle WebLogic Server installation for each group of domains and compute nodes that must be maintained independently.

To patch an Oracle WebLogic Server installation, you must use Smart Update. For more information, see *Oracle Smart Update Installing Patches and Maintenance Packs*.

9.8.4 Patching Software or Updating Firmware for Exalogic Machine Hardware Components

Oracle recommends that you patch software or update firmware for the storage appliance, switches, and ILOM in the Exalogic environment on a system-wide basis.

10

Monitoring the Topology Using Oracle Enterprise Manager Grid Control

Oracle Enterprise Manager Grid Control 11*g* with Oracle WebLogic Server Management Pack Enterprise Edition's capabilities include Exalogic-specific management tools to monitor Oracle software deployed in the Exalogic environment.

These monitoring capabilities expand on the existing Oracle WebLogic Server management features that span the following:

- Application performance management
- Configuration management
- Service-level management and operations

New Features

These features improve the performance and availability of Java applications and web services, avoid downtime, and reduce cost by automating manual and error-prone operations.

The following are the new Exalogic-specific features offered by Oracle Enterprise Manager Grid Control:

- View an Exalogic dashboard that shows the overall health, availability, and performance of all applications, Oracle WebLogic domains, and hosts running on Oracle Exalogic machine.
- Drill down into application deployments to identify metrics and set thresholds.
- Display alerts, policy violations, and incidents hierarchically across the Exalogic environment and provide operations with instant notifications in relation to service levels and thresholds.
- Contextually drill down from the Exalogic dashboard and underlying menus into the detailed component and JVM-level performance metrics, configuration management, and provisioning and cloning that provide end-to-end management for operations and administrators.

This chapter includes the following topics:

10.1 Accessing Oracle Enterprise Manager Grid Control 11*g*

You can access Oracle Enterprise Manager Grid Control 11*g* by navigating to the following URL:

```
http://<hostname.domain>:<port>/em
```

For example:

```
http://exalogic.mycompany.com:1159/em
```

The Resource Center is your central access point to Enterprise Manager documentation as well as the comprehensive technical resources of the Oracle Technology Network (OTN).

10.2 Discovering an Oracle Exalogic Target

You must discover the system targets for your Oracle Exalogic in the Enterprise Manager. To do so, complete the following steps:

1. Log in to Oracle Enterprise Manager Grid Control web interface. The home page is displayed.
2. On the home page, click the **Targets** tab, and then **Systems**.

   The **Systems** page is displayed.
3. Scroll-down and select **Exalogic Elastic Cloud** next to **Add**.
4. Click **Go**.

   The **Discover Exalogic Elastic Cloud** page is displayed.
5. Enter the following details:
   - Enter a unique name for the Oracle Exalogic target you want to monitor in the **Name** field. For example: **Exalogic Enterprise Deployment**
   - Select your **Exalogic ID** for the Oracle Exalogic machine you want to monitor. This ID is generated when you install Oracle Exalogic Machine.
6. Click **Finish**.

   A **Confirmation page** is displayed (Figure 10–1).

Figure 10–1 Confirmation Page

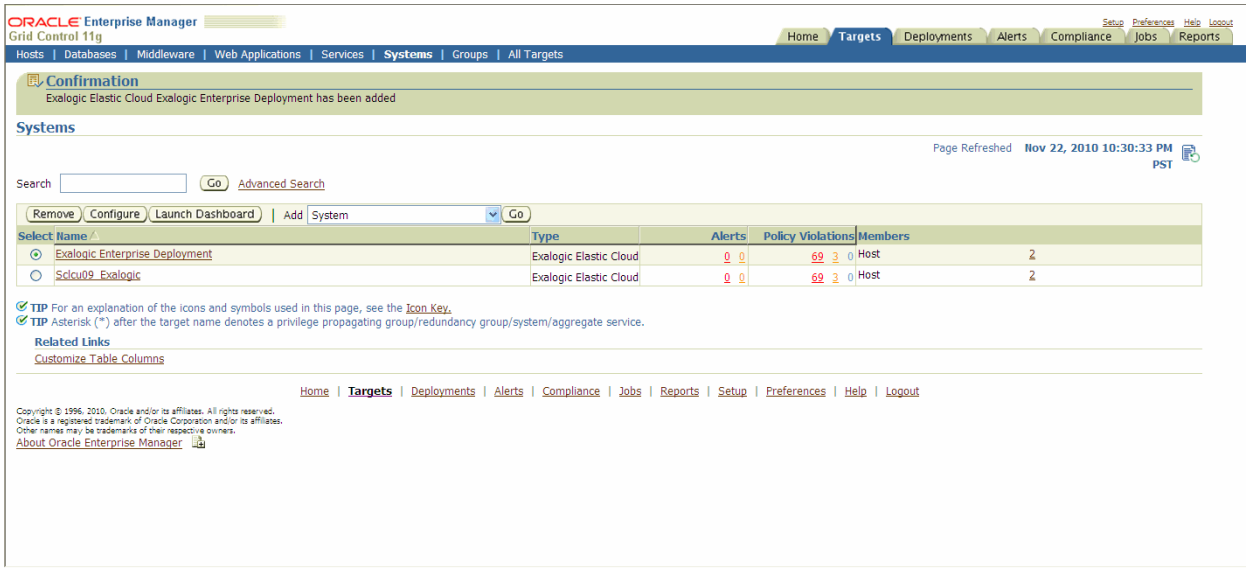

10.3 Using Exalogic-Specific Pages in Oracle Enterprise Manager Grid Control 11*g*

To navigate to the Exalogic-specific pages in Oracle Enterprise Manager Grid Control 11*g*, do the following:

1. Log in to Oracle Enterprise Manager Grid Control web interface. The home page is displayed.
2. On the home page, click the **Targets** tab, and then **Systems**.

   The **Systems** page is displayed.
3. On the **Systems** page, select the target name for your Oracle Exalogic machine, such as `Exalogic Enterprise Deployment`. This is the Target you have created in Section 10.2, "Discovering an Oracle Exalogic Target".

   The Oracle Exalogic Home page is displayed (see, Figure 10–2). This is the landing page for all Exalogic-specific monitoring and control operations, including configurations, application deployments, WebLogic domains, and metrics pages.

Figure 10–2 Oracle Enterprise Manager - Oracle Exalogic Home Page

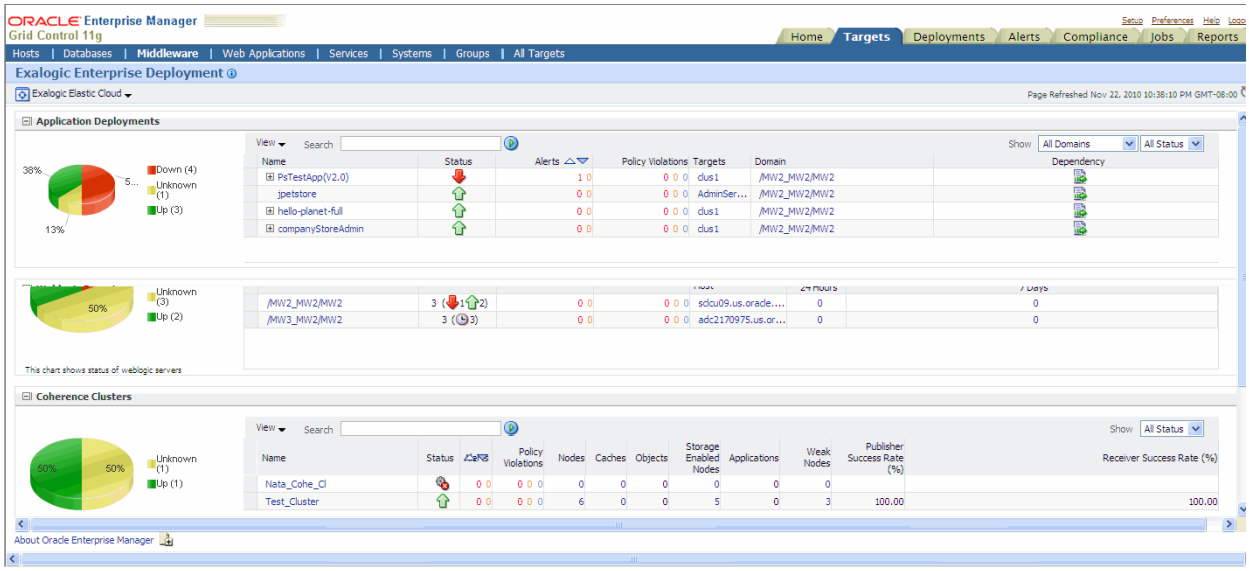

Oracle Enterprise Manager displays comprehensive metrics for Oracle Exalogic. When these metrics are displayed in tabular format, you can sort them in ascending or descending order. This feature enables you to find highest and lowest values easily.

The Oracle Exalogic home page, which is shown in Figure 10–2, displays various portlets. These portlets are described in the following sections:

- Section 10.3.1, "Application Deployments"
- Section 10.3.2, "WebLogic Domains"
- Section 10.3.3, "Coherence Clusters"
- Section 10.3.4, "Hosts"

10.3.1 Application Deployments

From **Exalogic Elastic Cloud**, click **Application Deployments**. The **Application Deployments** page is displayed.

Figure 10–3 Application Deployments

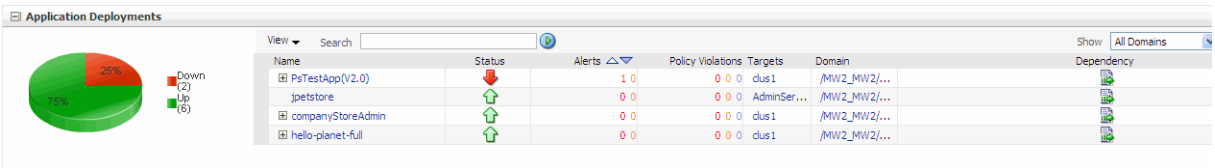

Use the Application Deployments page to do the following:

- View general information
- View servlets, JSPS. and EJBs
- View Work Manager
- View alerts and policy violations

10.3.2 WebLogic Domains

From **Oracle Exalogic**, click **WebLogic Domains**.

The **Oracle WebLogic Domains** page is displayed (Figure 10–4).

Use the Oracle WebLogic Domains page to do the following:

- View information about the domain, such as status, alerts and policy violations, and configuration changes
- View response and load information of the top 5 servers by average response times
- View the resource usage of the top 5 servers by CPU usage percentage
- See useful performance data related to JMS/JDBC/EJBs/JSPS and servlets of all servers across all domains

Figure 10–4 WebLogic Domains

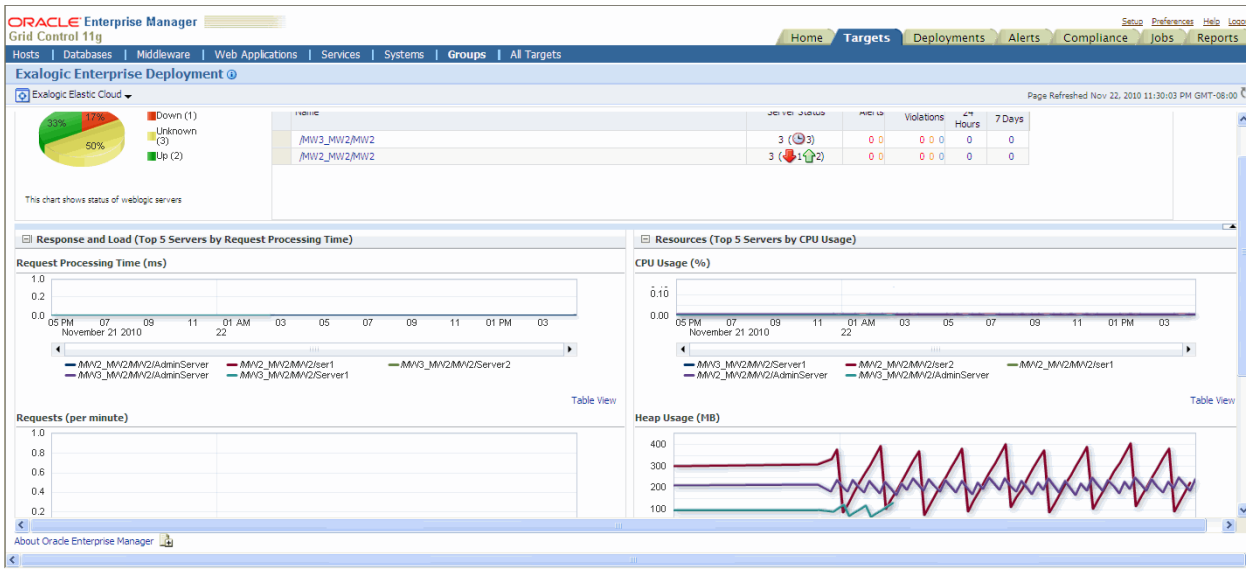

10.3.3 Coherence Clusters

From **Oracle Exalogic**, click **Coherence Clusters**. The **Coherence Clusters** page is displayed (Figure 10–5).

Use the Coherence Clusters page to do the following:

- View the Coherence Cluster name
- View status, alerts, and policy violation
- View number of nodes, caches, objects, applications, and weak nodes
- View Storage Enabled Nodes
- View publish success Rate and receiver success rate

Figure 10–5 Coherence Clusters Page

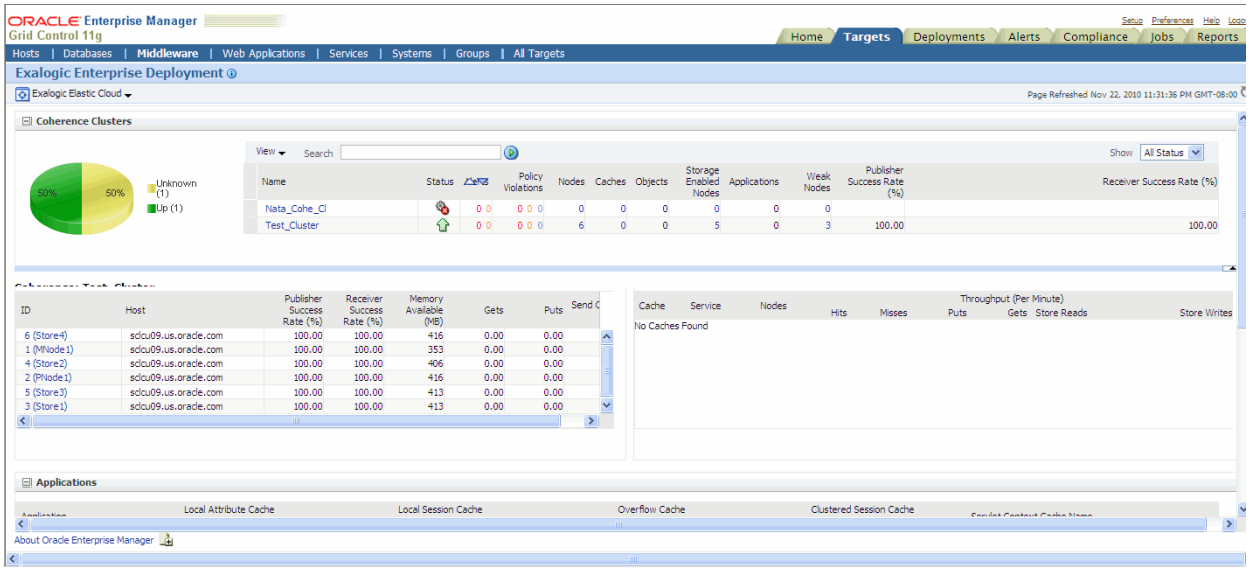

10.3.4 Hosts

From **Oracle Exalogic**, click **Hosts**. The **Host** page is displayed.

The Host Home page provides a glimpse of the vital statistics for this host, which is part of the greater Enterprise Manager environment.

In the Host page, select the Host you want to monitor from the table. Details specific to the host you selected is displayed (Figure 10–6). Using the Host Home page, you can:

- Drill down to view detailed statistics about this host

  For example, by clicking the **Local File Systems** metric in the Configuration section, you can see a summary of the local file system space being used. On the ensuing page, when you click the percentage associated with a mount point, the Metric Detail page for that mount point appears.

- Study the policy violations for the host
- Study all the alerts associated with this host
- Analyze the job activity
- Determine whether there are outstanding patch advisories
- Determine the last security evaluation of the host
- Investigate further the health of the host

Figure 10–6 Hosts Page

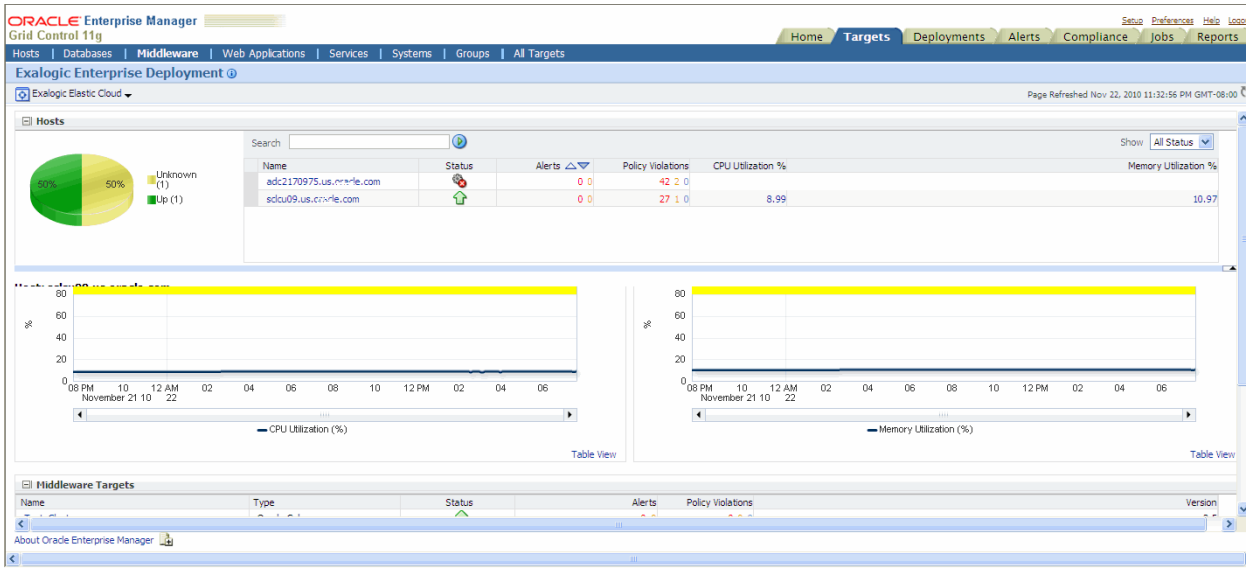

A

Migration Considerations for Oracle Exalogic

This chapter provides general guidelines and deployment configuration recommendations for migrating your existing application deployments to the Oracle Exalogic environment.

Your existing environments may vary based on the following:

- Number of applications in the domain
- Number of clusters and size of the clusters
- Number of servers in the domain

If you are migrating multiple domains, each of which may be different from each other. Given the differences in existing deployments, there is no set series of steps to follow when migrating to the Oracle Exalogic environment.

The chapter includes the following sections:

Migration Process

To migrate your existing environment, Oracle recommends the following steps:

- Review the domain configuration and apply the specific application deployment configuration guidelines. For more information, see Section 1, "Enterprise Deployment Overview".
- Determine the domains that will be migrated to the Oracle Exalogic environment.
- Determine the High Availability requirements for each application in the domain.
- Define the topology for the new Oracle Exalogic environment based on each domain's CPU and High Availability requirements. For more information, see Section 5, "Configuring Oracle Fusion Middleware".
- Migrate each domain, updating the domain configuration to reflect the new file system, recommended topology guidelines, and application deployment configuration changes.

Recommended Topology

Oracle recommends the following topology for the Exalogic environments:

- A domain spans at least two compute nodes.
- Each cluster spans at least two compute nodes.
- Each Managed Server is located on a server core.
- A node manager resides on each compute node.
- Domains, applications, and binaries are located on shared storage.

For more information, see Section 5, "Configuring Oracle Fusion Middleware".

Application Deployment Configuration Guidelines

In general, applications can be migrated to the Exalogic environment with minimal changes. Consider the following:

- Section , "Shared File System"
- Section , "Staging Model"

Shared File System

The Oracle Exalogic environment contains a shared file system available to all servers. When migrating your application, Oracle recommends that the application archives and domain directories be located on the shared file system. In an Oracle WebLogic Server domain, each individual server can either share a domain directory or have its own domain directory. In the Exalogic environment, Oracle recommends that one domain directory be shared among all servers.

Staging Model

Oracle WebLogic Server supports three types of deployment staging modes for deployed applications. For more information, see section "Controlling Deployment File Copying with Staging Modes" in the *Oracle Fusion Middleware Deploying Applications to Oracle WebLogic Server*.

In the Oracle Exalogic environment, application archives are located on the shared storage and one physical copy of the application is available to all servers. Because of this, Oracle recommends using the **nostage** staging mode for applications.